DISTRIBUTED SNIFFER SYSTEM™

Analyzer Operations Manual

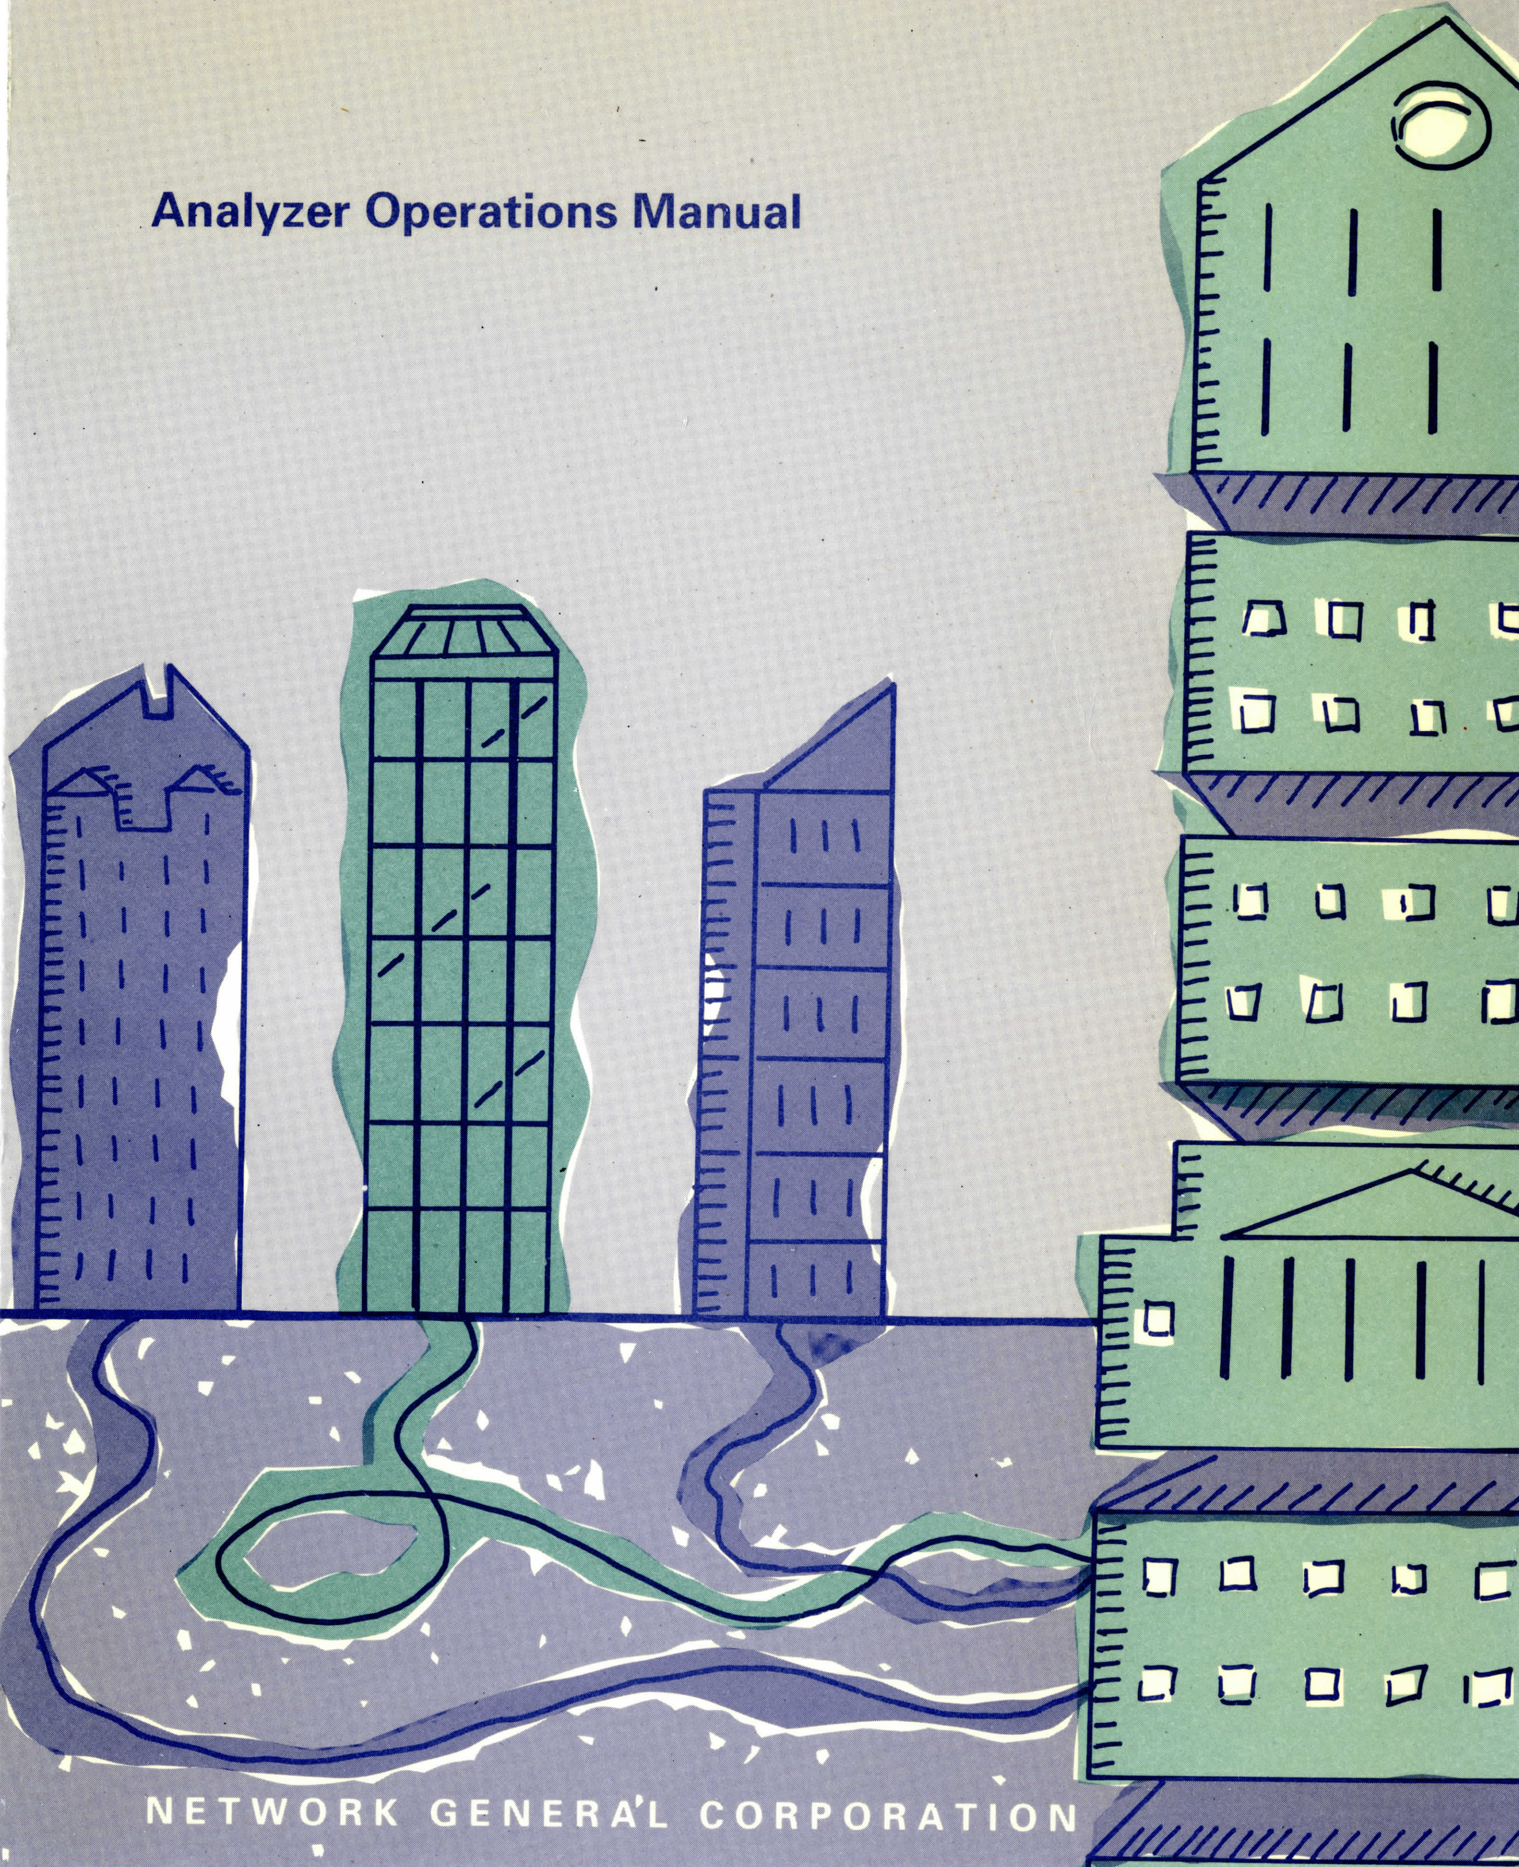

NETWORK GENERAL CORPORATION

DISTRIBUTED SNIFFER SYSTEM™

Analyzer Operations Manual

NETWORK GENERAL CORPORATION

DISCLAIMER OF WARRANTIES

The information in this document has been reviewed and is believed to be reliable; nevertheless, Network General Corporation makes no warranties, either expressed or implied, with respect to this manual or with respect to the software and hardware described in this manual, its quality, performance, merchantability, or fitness for any particular purpose. The entire risk as to its quality and performance is with the buyer. The software herein is transferred "AS IS."

Network General Corporation reserves the right to make changes to any products described herein to improve their function or design.

In no event will Network General Corporation be liable for direct, indirect, incidental or consequential damages at law or in equity resulting from any defect in the software, even if Network General Corporation has been advised of the possibility of such damages. Some states do not allow the exclusion or limitation of implied warranties or liability for incidental or consequential damages, so the above limitation or exclusion may not apply to you.

Document prepared by Paul Berry, with contributions from Len Shustek, David Trousdale, Florence Chan, and Terri FitzMaurice.
May 1991

P/N: 20037-001

Table of Contents

List of Figures

Preface

Chapter 1. Overview— What a Sniffer Analysis Server Does

Chapter 2. Monitoring Traffic and Capturing Frames

Chapter 3. Network-Specific Testing

Chapter 4. Generating Traffic to Load the Network

Chapter 5. Displaying and Interpreting Captured Frames

Chapter 6. The Analysis Server's Use of Files

Chapter 7. Protocol Interpreter Combinations

List of Procedures

Preface

Chapter 1. Overview— What a Sniffer Analysis Server Does

Chapter 2. Monitoring Traffic and Capturing Frames

Chapter 3. Network-Specific Testing

Chapter 4. Generating Traffic to Load the Network

Chapter 5. Displaying and Interpreting Captured Frames

Chapter 6. The Analysis Server's Use of Files

Chapter 7. Protocol Interpreter Combinations

DISTRIBUTED SNIFFER SYSTEM™

PREFACE

Network General

Preface

About This Manual

This manual describes the functions and operations of the Sniffer analyzer, a software component of the Distributed Sniffer System™.

The Distributed Sniffer System consists of two types of products: Sniffer® servers and SniffMaster™ consoles. Each server observes the local or wide-area network to which it is attached; the consoles control the servers and display the results of the servers' activities. Some servers run the monitoring or analysis application alone, while others run both.

Manuals for the Distributed Sniffer System

Two types of manuals accompany the Distributed Sniffer System. The primary manuals, which include this manual, describe the system's normal operations; the supplementary manuals describe the programs that configure and test the system's various hardware and software components for troubleshooting. Figure ii describes the primary manuals that accompany the Distributed Sniffer System. The actual manuals in your shipment depend on the configuration of your particular system.

For Information On...	Read...
Installing and configuring the server.	*Distributed Sniffer System: Installation and Operations Manual* or *Sniffer Server Installation Manual.*
Installing and configuring the console. Controlling the server from the console. Starting and terminating the applications on the server.	*Distributed Sniffer System: Installation and Operations Manual.*
Operating a server's *monitor* functions on a token ring network.	*Distributed Sniffer System: Token Ring Monitor Operations Manual.*
Operating a server's *monitor* functions on an Ethernet network	*Distributed Sniffer System: Ethernet Monitor Operations Manual.*
Operating a server's *analyzer* functions on Ethernet, token ring, or a WAN/synchronous network	*Distributed Sniffer System: Analyzer Operations Manual* (this manual)
Various network and protocol types.	*Distributed Sniffer System: Network and Protocol Reference.*

Figure i. Primary manuals for the Distributed Sniffer System.

In addition to the manuals that directly concern the Distributed Sniffer System, depending on the network or configuration you are using, you may also wish to consult the supplementary manuals listed in Figure ii.

For Information On...	Read...
Running the adapter diagnostics to test the IBM 16/4 token ring adapter in the console.	*Token-Ring Network Guide to Operations.*
Running the diagnostics to test the InterLan NI5210 Ethernet controller in the console.	*NI5210 Installation Manual.*
Configuring and using the IBM® Local Area Network (LAN) Support Program.	*Local Area Network Support Program, version 1.2, User's Guide.*

Figure ii. Supplementary manuals.

If the product shipment includes release notes or README files on disk, the information in the note or file supersedes the information in this manual.

Audience of This Manual

The manual has been prepared with the following assumptions:

- You are a network manager or troubleshooter who understands your network's operation.
- You are familiar with DOS.
- You have properly started the SniffMaster console.

Organization of This Manual

Figure iv describes this manual's organization.

Chapter	Contents
Table of Contents List of Figures List of Procedures Preface	For many tasks, the manual includes a step-by-step procedure; they're all listed in the *List of Procedures.*
Chapter 1, "Overview—What a Sniffer Analysis Server Does."	Provides an overview of the analyzer and describes its capabilities. It also discusses the menu structure.
Chapter 2, "Monitoring Traffic and Capturing Frames."	Describes the preparation required before you start to capture frames for analysis, and what the analyzer displays as capture proceeds.
Chapter 3, "Network-Specific Testing."	Describes the cable-test feature of the Ethernet analyzer.
Chapter 4, "Generating Traffic to Load the Network."	Describes the traffic generator for Ethernet and token ring networks.
Chapter 5, "Displaying and Interpreting Captured Frames."	Describes the procedures for filtering, interpreting, and displaying captured frames.
Chapter 6, "The Analysis Server's Use of Files."	Describes the server's directory structure and the types and formats of files it uses.
Chapter 7, "Protocol Interpreter Combinations."	Describes the procedure to generate Sniffer analyzers with alternative combinations of protocol interpreters.
Index	

Figure iv. Outline of this manual

Navigational Aids Used in This Manual

To help you find procedures easily, a separate list of procedures is provided in this manual in addition to the Table of Contents and List of Figures. To facilitate use of this manual as a reference, there is an extensive index.

This manual uses icons in the margin to help you locate important information as explained below:

The paragraph next to this icon contains information that is especially important; you should be certain to read it carefully before you proceed.

A warning gives you instructions that you must follow to avoid possible damage to data files, program files, or hardware devices.

A cautionary paragraph provides information that may help you avoid a possible pitfall or misunderstanding.

A recommendation describes a useful and valuable way of using the product.

A procedure is a series of steps for accomplishing a particular task.

Conventions Used in This Manual

Special Notations

The following describes the conventions used in this manual:

Bold — Menu options are in bold type. For example:

Move the highlight to **Display** and press Enter.

UPPERCASE — The filenames and command names you type at a DOS prompt are in uppercase. For example:

Names are recorded in the file STARTUP.*xx*D.

Bold italics — Variables, for which you insert values, are in bold italics. For example:

Type the number of minutes and seconds in ***mm:ss*** format.

`Screen font` — Screen messages are printed in monospaced font. For example:

If a monitoring session is in progress, the following message appears:

```
You must stop monitoring before you can use this feature.
```

Terminology

Hexadecimal numbers mentioned in the manual are followed by "(hex)"; numbers without any notations are decimal. For example, "The maximum number of stations is 75. The default memory address is D8000 (hex)."

The terms "Sniffer monitor" and "Sniffer analyzer" refer to software applications running of the Sniffer server. The term "SniffMaster console" refers to a software application running on dedicated PC, at which you enter all input and receive all output from the server.

Screen Displays and Keyboard Input

All the keystrokes mentioned in the manual are entered from the keyboard of the SniffMaster console. Similarly, all the screen displays generated by the monitor appear on the console's screen.

The screen displays in this manual may not be exactly the same as what you see on your console screen. For example, you can choose to have the console show the server name on each monitor display, but the screens in this manual do not generally show the name.

Other Sources of Information

Network General Corporation (NGC) provides other sources of information that can help you get familiar with the distributed Sniffer system.

On-Line Help

After highlighting an item in a console, analyzer, or monitor menu, you can see a phrase or sentence in a panel near the bottom of the screen. It explains the meaning of the highlighted item.

If you want to obtain general information on a particular feature of the Distributed Sniffer System, press F1 at any time. A dialog box containing a list of topics opens. If you are displaying a monitor statistics screen, pressing F1 gives you information on the current screen.

Other Manuals

This manual does not describe the characteristics of different protocol and network types. For general information on these topics, refer to *Distributed Sniffer System: Network and Protocol Reference.*

Tutorial

NGC distributes a booklet with accompanying diskette entitled *Real Networks. Real Problems.* It presents case studies based on data captured with a Sniffer network analyzer from four different networks. The Sniffer analyzer and the server's analysis application have different capabilities, but the case studies allow you to see how investigation of a network problem proceeds.

You can obtain the tutorial free of charge from any of the company's sales representatives or directly from NGC.

Technical Support

The procedure for obtaining technical support for problems with the Distributed Sniffer System is described in Appendix A of *Distributed Sniffer System: Installation and Operations Manual.*

CHAPTER ONE: OVERVIEW—WHAT A SNIFFER ANALYSIS SERVER DOES 1

Chapter 1. Overview—What a Sniffer Analysis Server Does

Chapter Overview

This chapter provides a general orientation to the Sniffer analysis server. It is aimed mainly at the reader who hasn't yet used a Sniffer analyzer. The chapter starts with a general description of the distributed analysis system and its four types of components: SniffMaster *consoles,* Sniffer *monitor servers,* Sniffer *analysis servers,* and Sniffer *combined monitoring and analysis servers* (combining the functions of monitoring and analysis in a single server).

Each server has two channels, one for *monitoring* the network being studied, the other for *transporting* data and commands between the server and the console. The server is a passive observer of a local area network (LAN) or of the synchronous link of a wide area network (WAN).

The chapter depicts the organization of a server's *analysis* functions by a map of the way information flows between them and back to the console. Analysis has two main phases: *capture* (when frames are recorded in a storage buffer), followed by *analysis* (when the captured frames are interpreted and displayed). Capture can occur *live* from the network, or as a *replay* of previously captured data.

Before capture starts you can set *filters* to avoid filling the buffer with extraneous traffic. You can set a *trigger pattern* that causes the server to stop capture when it has found the condition you specify. While capture proceeds, you have a choice of tabular or graphic displays that indicate the density of traffic.

Analysis follows capture. The captured frames are displayed in three formats: *summary, detail,* and *hex.* Frames and their embedded messages are interpreted at each level. Addresses within the frames are interpreted from a table of *symbolic names.* You can filter the display, or search through it.

Because protocol interpreters are so numerous and so large, it's impossible (and unnecessary) to install all of them at once. Each analyzer is initially built with two combinations of protocol interpreters. The *configuration utility* lets you construct other combinations as you wish.

The Distributed Sniffer System

The Sniffer analysis server whose operations are described in this manual is one of the three principal components of the Distributed Sniffer System. For a perspective on the entire system and the role of the SniffMaster console in controlling it, see the companion publication, *Distributed Sniffer System: Installation and Operations Manual.*

The Role of the Analyzer

An analyzer —that is, the network analyzer program installed in a Sniffer analysis server— records and interprets network transmissions. The work of analysis occurs in two main stages:

Capture The analyzer records network traffic for later interpretation. Capture can be filtered to record only traffic meeting certain criteria. Capture can be frozen when a triggering condition is observed. This assures that the retained sample includes traffic just before or after the event of interest.

While capturing frames, the analyzer software maintains graphs or tables that summarize the traffic it has recorded.[1]

Display The analyzer interprets the recorded traffic. During display, the analyzer decodes the various layers of protocol in the recorded frames and displays them as English abbreviations or summaries. The analyzer can filter the display to show only those frames that meet certain criteria.

The Analyzer as (Mostly) Passive Observer

A Sniffer server "hears" all traffic that passes through the segment it is monitoring. On a WAN/synchronous link, it hears traffic in both directions ("from DTE" and "from DCE"). On a LAN, it hears all traffic that passes through the segment or subnet that it is monitoring. It is characteristic of a LAN that every station physically receives every transmission. Ordinarily, each ignores all messages except those addressed to it. The Sniffer server not only hears all transmissions but, while in "capture" mode, can record them, regardless of how they're addressed.

In general, the Sniffer server observes, tabulates, or captures, but contributes nothing of its own to traffic on the network it is monitoring. However, when a server is monitoring a LAN, the server may contribute to the LAN's traffic as follows:

- On Ethernet or token ring, an analysis server can generate test frames. In that mode, it repeatedly transmits a single test packet you have specified.
- On Ethernet, the server can emit a pulse to test for cable defects.

[1] The analyzer's displays during capture resemble some of the displays produced by the monitor server. Don't confuse this mini-monitoring during capture with the full-blown monitor server, which is quite separate.

- On Ethernet or token ring, if the Sniffer server is configured so that its control link is to the same network segment that it is monitoring, the monitor channel can hear and capture the server's own control messages. In this circumstance, the server is taking part in the traffic that it is monitoring.
- On token ring, every station must participate in the ring by forwarding traffic from its upstream neighbor to its downstream neighbor. Each Sniffer server does that in the same way as other stations. However, the server does not reply to the poll for standby monitors, and never acts as the ring's active monitor. It is thus invisible to most other stations.
- On token ring, the server periodically transmits a frame addressed to "LAN Manager" announcing "trace tool present." The LAN Manager can force such a station to leave the ring immediately.
- On Ethernet or token ring, the server's transport card may be connected to the same network segment or ring as the server's monitor card. In that case, the server's monitor card observes the server's own transport activities. The transport activities rarely amount to more than a small percentage of network utilization. They use a higher-level protocol called NGCP. If you wish, you can set display filters to include or exclude them.

A Map of the Analyzer's Functions

The analyzer's activities are divided into the set of functions described below. The diagram in Figure 1–1 represents schematically the route by which information flows from between the various parts. Following the paths that frames may follow as they are captured from the network, the analyzer's principal components are as follows:

- Capture filter: it decides which arriving frames to discard and which to retain.
- Capture displays: these chart the progress of capture.
- Trigger detector: it scans arriving frames for a pattern. When it detects the pattern, it stops capture so that frames preceding or following the trigger event are retained.
- Capture buffer: this is a storage area for frames that have been accepted. From here they are subsequently interpreted and displayed.
- Protocol interpreters: these identify the protocols nested within each frame and interpret their contents.
- Display filters: these select from the capture buffer the frames that will be displayed.

- Display of the selected frames in three modes:
  - Summary
  - Detail
  - Hex with ASCII or EBCDIC
- Output of the display
  - Transmitted to the console screen
  - Saved to a disk file on the server
  - Sent to a printer

Figure 1–1 represents the route that frames follow as they move from the network to the network interface card, from there through filters and triggers to the capture buffer, and from there to the interpreters and displays. You may want to turn back to the figure as you read descriptions of these functions in the sections that follow.

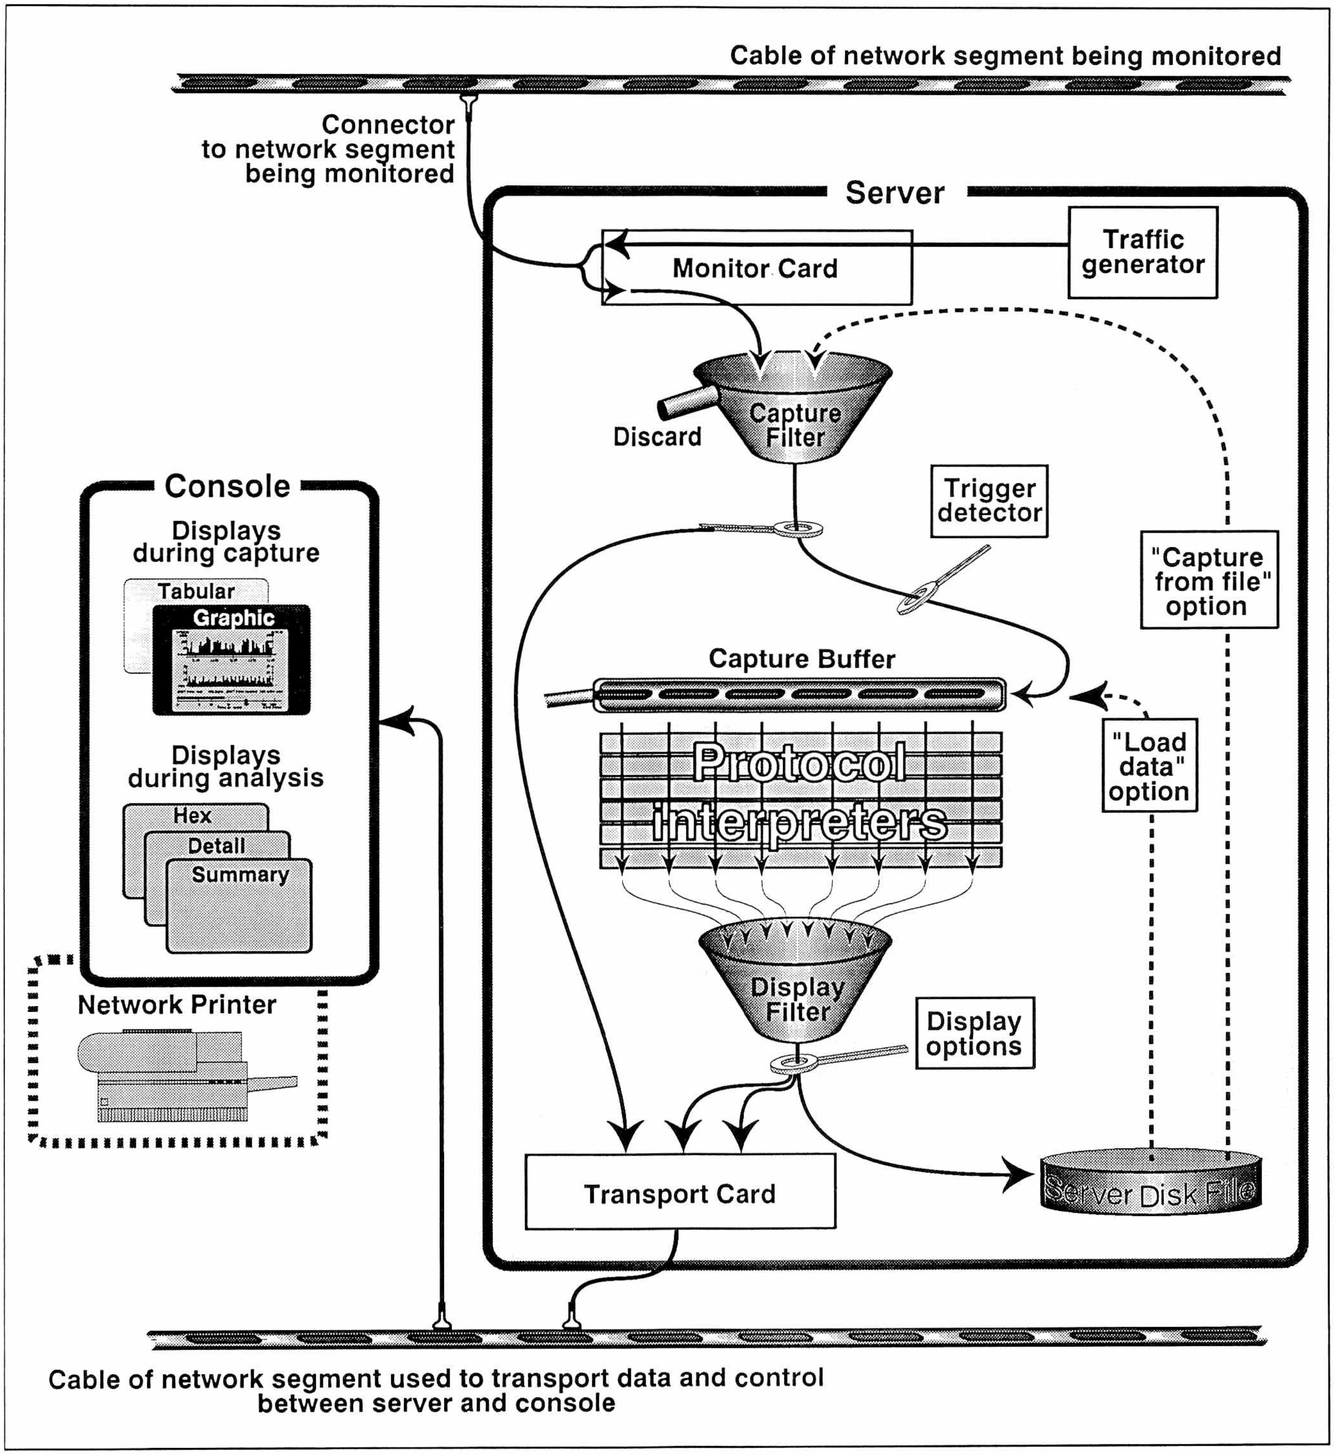

Figure 1–1. Schematic view of the flow of frames in the Sniffer analyzer.

Capturing Traffic

The analysis server scrutinizes the traffic it hears and records the frames that pass various filters. Recorded frames may be saved to a file for analysis later, or immediately displayed, filtered, and interpreted.

Sources of Capture— Live vs. Recorded

Before the Sniffer analyzer can provide detailed interpretation or analysis, it must obtain a sample of network traffic. There are two ways to do this:

- Live capture, or
- Capture from a previously recorded file

Live capture

The Sniffer analysis server's network interface card detects all traffic, and passes it to the analysis server's CPU. Frames that pass the capture filter are stored in the capture buffer for subsequent display or analysis. As frames arrive, the analysis server updates its real-time display of traffic density.

Capture from a previously recorded file

Instead of capturing from the network, the Sniffer analyzer can read from a file of saved frames. Such a file is created by saving the contents of the capture buffer. By capturing from a file, you can "replay" a capture that took place at an earlier time. Since a file of captured frames includes a timestamp for each frame, the replayed capture reproduces not only the data but also the timing of each frame's arrival.

When you're capturing from a previously recorded file, you can set capture filters just as you would for live capture, accepting into the capture buffer only those frames that pass the filter. Capturing from a file is a way to gain experience in operating the analysis server or setting its capture filters. You can replay recorded captures even when the network that the server monitors is not active or the server's "monitor" card is not connected.

Filters to Select Frames for Capture

The number of frames reaching the adapter is potentially so large that it's often essential to select only a subset for storage. Saving them all would take far too much space. Before you start capture, you can set filters so that the analysis server saves only frames that meet certain criteria. The analysis server applies a filter to each newly-arrived frame and discards all frames that do not meet the test.

Capture filters may be any combination of the following types. In each case, the filter can be set to save frames that meet the test or those that don't.

Station address
The Sniffer analyzer accepts frames sent to or from particular addresses that you specify.[1]

Destination class
The Sniffer analyzer accepts frames that have specific destinations, rather than frames addressed to a group (for example, broadcast).

Unknown station
The Sniffer analyzer accepts frames sent from or to an "unknown" station— that is, a station not included in the table of names for stations.

Protocol
The analyzer accepts frames that contain any of the low-level protocols you specify. (During capture, to maintain speed, filtering is restricted to the lower levels. During display, you can set filters for higher level protocols embedded within the frames.)

Pattern
The Sniffer analyzer accepts frames that contain a specified pattern. A pattern can be constructed as some logical combination of up to eight component patterns. Each of the component patterns is described by a particular sequence of bits or characters at specified positions. For example, a filter might accept only frames that pass between a particular user and server and involve a particular error code in a particular protocol.

Displays During Capture

While it is capturing frames, the Sniffer analyzer is able to tabulate and meter the frames it is capturing. The tabulations in some ways resemble those provided by the Sniffer monitor, a separate software module available for Ethernet and token ring, routinely installed on each analysis server for those networks. However, the Sniffer monitor doesn't save frames for later analysis and doesn't interpret their protocols. When your interest is primarily monitoring, you may find it preferable to run the Sniffer monitor rather than the Sniffer analyzer. You invoke the monitor functions from the main selection menu (see Figure 1–6, page 1–22).

[1] During capture, you can only filter on DLC addresses. During display, you can filter both by address level or on specific high-level addresses.

What follows describes the more limited facilities for metering and tabulation that accompany the capture phase of the Sniffer analyzer.

Before you start capture, select the type of display the Sniffer analyzer will be generating while frames arrive. The choices are:

Graphic "Skyline" displays of traffic density over time.

Tabular Source or source-and-destination counts in various formats.

Highspeed On Ethernet, this option suppresses displays during capture to let the analyzer give full attention to arriving traffic.

Traffic Density Displays

While it "listens" to network traffic, the analysis server accumulates statistics. As frames arrive, it either displays real-time graphs of traffic density, or updates a table that counts messages by source and destination. An example of a traffic density graph is shown in Figure 1–2. A new column is added at the right once every second, minute, or hour, depending on the time scale you select.

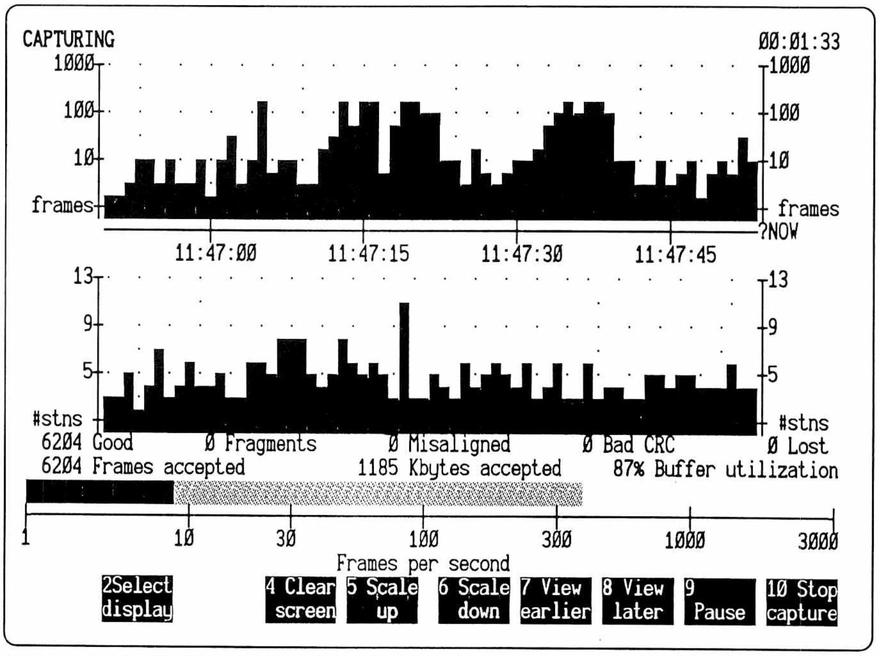

```
CAPTURING                                                  00:01:33
6204 Good       0 Fragments       0 Misaligned       0 Bad CRC      0 Lost
6204 Frames accepted          1185 Kbytes accepted       87% Buffer utilization
```

Figure 1–2. Skyline view of traffic density at one-second intervals.

The display shows two histograms, one above another. On a synchronous link, the display shows histograms for traffic from DTE and from DCE. On a LAN, one shows traffic density (frames or kilobytes per second) and a second the number of stations that are active (have transmitted during the time interval).

Source and Destination Tables During Capture

As an alternative to the "skyline" histograms, the analyzer can instead tabulate the frames as it accepts them, with separate counters for each source-and-destination pair.

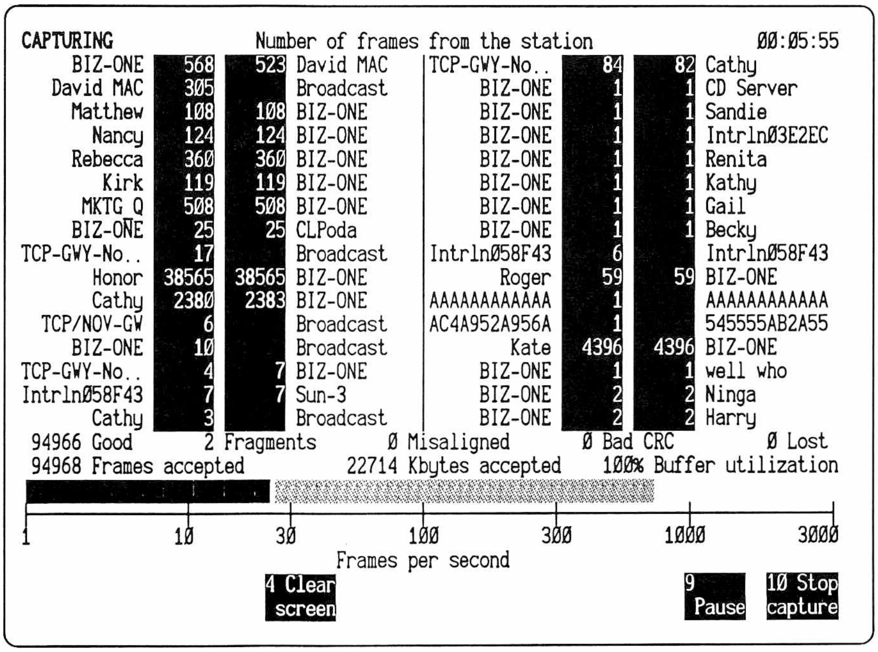

```
CAPTURING                Number of frames from the station                 ØØ:Ø5:55
      BIZ-ONE   568    523 David MAC   |TCP-GWY-No..   84    82 Cathy
    David MAC   3Ø5        Broadcast   |     BIZ-ONE    1     1 CD Server
      Matthew   1Ø8    1Ø8 BIZ-ONE     |     BIZ-ONE    1     1 Sandie
        Nancy   124    124 BIZ-ONE     |     BIZ-ONE    1     1 IntrlnØ3E2EC
      Rebecca   36Ø    36Ø BIZ-ONE     |     BIZ-ONE    1     1 Renita
         Kirk   119    119 BIZ-ONE     |     BIZ-ONE    1     1 Kathy
       MKTG Q   5Ø8    5Ø8 BIZ-ONE     |     BIZ-ONE    1     1 Gail
      BIZ-OÑE    25     25 CLPoda      |     BIZ-ONE    1     1 Becky
TCP-GWY-No..     17        Broadcast   |IntrlnØ58F43    6       IntrlnØ58F43
        Honor 38565  38565 BIZ-ONE     |       Roger   59    59 BIZ-ONE
        Cathy  238Ø   2383 BIZ-ONE     |AAAAAAAAAAAA    1       AAAAAAAAAAAA
  TCP/NOV-GW      6        Broadcast   |AC4A952A956A    1       545555AB2A55
      BIZ-ONE    1Ø        Broadcast   |        Kate 4396  4396 BIZ-ONE
TCP-GWY-No..      4      7 BIZ-ONE     |     BIZ-ONE    1     1 well who
IntrlnØ58F43      7      7 Sun-3       |     BIZ-ONE    2     2 Ninga
        Cathy     3        Broadcast   |     BIZ-ONE    2     2 Harry
 94966 Good          2 Fragments          Ø Misaligned      Ø Bad CRC          Ø Lost
 94968 Frames accepted          22714 Kbytes accepted      1ØØ% Buffer utilization
```

Figure 1–3. Source-and-destination pair count during capture.

During capture, traffic on a LAN is subtotaled by station, either individually by source, or pairwise by source and destination. Figure 1–2 shows a tabulation by source and destination. On a synchronous link, counts are subtotaled by type of frame and also by higher-level address (logical call number). There are separate counts for traffic from DTE and from DCE.

The Capture Buffer

Frames that are accepted move to the capture buffer. The capture buffer is a large block of the server's main memory. The exact amount varies with the amount of RAM installed in the server and the amount reserved by the server and analyzer software. The capacity ranges from a few thousand medium-sized frames, up to many thousands. Where space is at a premium or frames are long, you can save space by truncating frames that exceed a given length.

Frames accumulate in the capture buffer in the order they are received. When the capture buffer becomes full, the Sniffer analyzer can halt capture or it can discard older frames to make way for new arrivals.

All analysis and interpretation is done on frames in the capture buffer. While the buffer has substantial space —room for thousands of captured frames— it is still very small compared to the volume of traffic on a network. That's why it's essential to restrict the buffer to frames that are relevant. There are two main ways to do this:

- Accepting frames selectively, so that the buffer isn't flooded with irrelevant frames. This is done by the *capture filters* command (described in the preceding sections).
- Timing the end of capture so that the frames that interest you are still in the buffer. This is done by the *trigger pattern detector* (described in the next section).

Stopping Capture When a "Trigger" Pattern Occurs

To be sure that the frames that interest you are retained in the capture buffer, you can set a detector that freezes capture when it sees a frame containing a pattern you've specified. The detector scans the stream of incoming frames *after* they've passed the capture filters but *before* they go to the capture buffer. (See Figure 1–1, page 1–7).

When it finds the trigger pattern, the detector stops capture. You can examine the frames in the buffer to see what has happened. You can have the detector stop capture immediately or after some delay. Setting a delay lets the server retain frames that follow the trigger event, as well as frames that precede it.

A trigger pattern is some combination of bits, bytes or characters at various positions in a frame. A trigger can be constructed as a logical combination of up to eight component patterns, each described by a particular sequence at specified positions.

To illustrate use of a trigger pattern, suppose it appears that there are intermittent difficulties with access to a file server. A filter could restrict capture to frames to or from the server. A trigger pattern could be set to halt capture when a frame from the server contains an error code in a protocol related to file transfer. When the analysis server reports such a frame, you have both a record of the error report and of the sequence of events that led to it.

Alternatives When Capture Is Complete

Once capture has stopped, you can:

- **Copy** the contents of the capture buffer to a file for later analysis or display.
- **Filter** the display, so you see only those frames that meet certain criteria.
- **View** the frames through one or more views, each with a particular style of display.
- **Browse** through the frames in the capture buffer.

- **Print** the contents of the buffer, according to the filters and views you've specified.

You have many options for displaying the contents of the capture buffer, either on the SniffMaster screen or in a printed report. (You can direct output either to a printer or to a file. The printer may be attached to the server or to the console; the file must be on the Sniffer server's hard disk (but can then be uploaded to the SniffMaster console if you wish).

You can set up a display filter so that frames that don't interest you are omitted from the display (even though they remain in the capture buffer). The mechanism for filtering frames from the display is like the mechanism for filtering frames during capture.

Saving the Capture Buffer

From the keyboard, you can select a command that saves the contents of the capture buffer to a file. You can save the entire capture buffer or just the frames that are accepted by your current display filter.

Once you have saved the capture buffer to a file, you can:

- Load the capture buffer with the saved file.
- Capture from the saved file— that is, re-enact the original capture.

Either procedure restores the capture buffer to the way it was when you saved it (but without frames that you excluded at the time you saved).

Interpreting What You've Captured

The Sniffer analyzer takes what would otherwise be a dense and meaningless stream of bits and translates it into readable English. It dissects each transmission (called a frame or packet) into its component layers. Then it decodes each layer according to its protocol—the rules of format and syntax that govern its encoding. For a quick summary view, the analysis server can show you just the highest layer of interpretation for each embedded packet. Or, if you request, it can show you all layers from lowest to highest. For each layer, the interpreter can show a succinct one-line summary, or a detailed translation of the important fields and parameters.

Protocol Interpreter Suites

The Sniffer analyzer doesn't just capture and store frames from the network. It also interprets them. When you select the detail view, you get a set of interpretations for each frame, one interpretation for each level of protocol that the frame contains.

Interpretation of the lowest-level protocols is provided automatically when you specify the network to which your Sniffer analysis server will be attached. Separate protocol interpreter suites interpret higher-level protocols. Each suite provides interpretation of several protocols likely to occur together. Figure 1–5 shows a list of the protocol interpreter suites available when this manual was printed; it's on page 1–21 (with the discussion of alternate configurations).

Each analysis server has all of Network General's current protocol interpreter suites. However, only a subset is enabled at one time (see "Configuring Alternative Sets of Protocol Interpreters" on page 1–20).

Symbolic Names for Addresses

To make its displays easier to read, the interpreter augments the binary codes that identify sender, receiver, or routing with symbolic entries from a name table. Thus, when a packet's high-level destination is 00100100001101010000000011000011, rather than showing the address as 24 35 00 C3 or even [36.53.0.195], the analysis server lets you see it as "Janet's Workstation" or whatever name makes sense to you. When the analysis server has no name for a station, it displays the name of the manufacturer and the unique address that the manufacturer assigns to that station.

Names and Name Tables

A *name table* pairs a numeric network address with an arbitrary name. When the Sniffer analyzer identifies a station, it can substitute the name in the table for the numeric code that the server actually received.

The first time you display frames after a new capture, the analysis server checks through all of the frames in the capture buffer and adds to its working copy of the name table any DLC addresses that are not yet in it. You can edit the working table to provide symbolic equivalents for the addresses that lack them. You can also save the working name table and reuse it during subsequent captures or displays.

After you have provided symbolic equivalents, the analysis server uses them not only in displays but also during subsequent capture sessions. The Sniffer monitor shares tables with the Sniffer analyzer, so that they both benefit from additions to the name table.

Resolving Names from Other Tables

The **manage names** menu contains an instruction to *resolve names*. When you select this option, the analysis server searches a saved name table to find equivalents for addresses that are not yet named in the analyzer's working name table.

Searching for Names

In some network protocols (for example, NetBIOS, AppleTalk or Novell NetWare), stations may declare names for themselves. The Sniffer analyzer is able to search the capture buffer for frames that contain names and use the names it finds to fill blanks in the working name table.

Saving and Restoring Setups

The Sniffer network analyzer offers a rich choice of options concerning both capture and display. To simplify work at subsequent sessions, you can name and store a record of the options you've selected. That's called a *setup* file. At a subsequent session, you have only to load a previously-saved setup, and all the settings it contains are restored.

Three Ways to View the Captured Frames

The Sniffer analyzer offers three levels of interpretation:

Summary	A one-line summary of the highest level of protocol in each frame, or, if you wish, of each of the various levels each frame contains.
Detail	A detailed English translation of all fields and parameters within the frame.
Hexadecimal	Exact record of the transmitted data, with ASCII or EBCDIC transliteration, as appropriate.

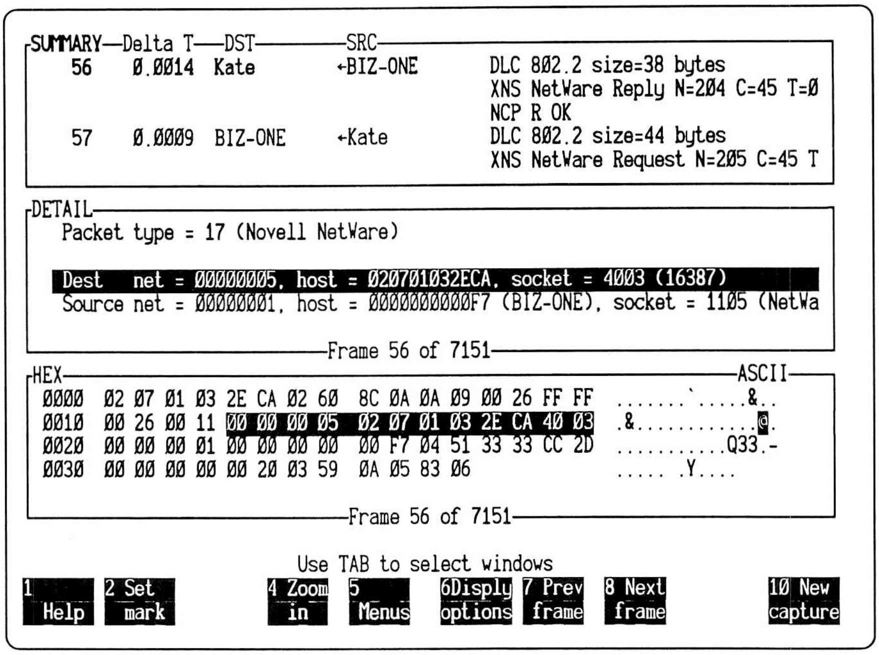

```
SUMMARY—Delta T——DST————————SRC————————————————————————————————————————
  56    0.0014  Kate         ←BIZ-ONE        DLC 802.2 size=38 bytes
                                             XNS NetWare Reply N=204 C=45 T=0
                                             NCP R OK
  57    0.0009  BIZ-ONE      ←Kate           DLC 802.2 size=44 bytes
                                             XNS NetWare Request N=205 C=45 T

DETAIL——————————————————————————————————————————————————————————————————
  Packet type = 17 (Novell NetWare)

  Dest   net = 00000005, host = 0207010032ECA, socket = 4003 (16387)
  Source net = 00000001, host = 000000000F7 (BIZ-ONE), socket = 1105 (NetWa
——————————————————————————————Frame 56 of 7151——————————————————————————
HEX—————————————————————————————————————————————————————————————ASCII———
 0000  02 07 01 03 2E CA 02 60  8C 0A 0A 09 00 26 FF FF   .......`.....&..
 0010  00 26 00 11 00 00 00 05  02 07 01 03 2E CA 40 03   .&............@.
 0020  00 00 00 01 00 00 00 00  00 F7 04 51 33 33 CC 2D   ...........Q33.-
 0030  00 00 00 00 00 20 03 59  0A 05 83 06              ..... .Y....
————————————————————————————————Frame 56 of 7151————————————————————————

                         Use TAB to select windows
```

Figure 1–4. Three views in display of a captured frame.

In Figure 1–4, all three views are visible: summary, detail, and hex.

Detail view

Each frame is decoded to show the type of frame and the protocols it contains. For high-level frames, the interpretation may require several levels, one for each protocol the frame contains. For each protocol, the detail view names the principal fields and interprets the values of each.

For each address in each protocol, the detail view shows both the numeric code actually transmitted and a symbolic equivalent. You can edit the address table to insert your own names, or import names from an external file.

Summary view

This condensed form abbreviates and truncates some of the information from the detail view.

All-level vs. highest-level-only display

You can elect to display a one-line summary of each protocol within a frame, or to show only the highest level.

Hexadecimal view

The entire frame is listed. You can elect to have character data displayed according to ASCII or EBCDIC conventions.

The hexadecimal view and the detail view show data for just one frame. The summary view shows not only the frame you are examining but a few on either side of it as well for context.

Distinguishing Protocol Layers

The interpreter labels and decodes the standard fields in each frame, making it easy to see the message conveyed. At the SniffMaster console, each layer is shown in a distinctive color (except, of course, when you equip the console with a monochrome display). Color coding applies in all three views: summary, detail, and hex.

Coordination of views

When you open two or more views at once, the Sniffer analyzer coordinates highlighting in the various views. The detail view's interpretation is automatically scrolled to match the level you've highlighted in the summary view. When both the detail and hex views are open, moving the highlight in the detail view produces a corresponding highlight in the hex view, so you can see which parts of the hex display correspond to the interpreted detail.

Two-Station Format

Frequently, analysis concerns the flow of commands between a pair of stations. Two-station format is a variant form of the summary view. Frames from the first station are summarized on the left side of the screen, and frames from the second station are summarized on the right. That way, it's easy to distinguish the two sides of a conversation. (If you also accept frames from other stations, they remain in the normal full-width format.)

Views and Viewports

Each view you select appears in its own rectangle within the displayed screen. The screen may be tiled as one, two, three, four, or six equal-sized rectangles, according to the number of views and viewports you request.

Two viewports

Sometimes it's important to compare a frame from one part of the capture buffer with a frame that arrived earlier or later. You can do that by electing two viewports. That splits the screen into left and right halves. You can scroll the two sides independently, permitting you to display one frame on the left and another on the right.

Scroll

When a frame's display doesn't fit within its panel, you can scroll the active panel both vertically and horizontally until the part you want is in view.

Zoom

You can temporarily assign one view to the entire Sniffer display. To do that, press F4, labeled zoom. The active view gets the entire display area until you again press F4 to "zoom out," thereby restoring the other views.

Filtering and Searching During Display

During display, you can use filters to select frames from the capture buffer. They're called display filters; they work in much the same way as capture filters. However, they can also respond to criteria that are not feasible during capture. For example, they can choose to display only frames that contain a particular high-level protocol or high-level address.

Filtering permits you to suppress traffic that's irrelevant to your concern, and to limit the display to frames involved in a particular transaction. With the obscuring bulk cleared away and the remaining frames concisely interpreted, it becomes easy to trace interactions between stations. You can see normal patterns at a glance, and readily identify suspicious exceptions.

It's also possible to search through the captured frames for those that contain particular items. This feature not only permits you to search for a particular address or pattern of data, but also for text that the Sniffer analyzer uses in the interpretation. Thus, you can search for names, values, or phrases. Frequently, the text you search for is implied by codes contained in the frame, but present as characters only in the Sniffer analyzer's interpretation.

Low and High-Level Filtering

Display filters work in much the same way as capture filters. However, because display filters don't have to work under time-pressure, they can also examine higher-level protocols and higher-level addressing. The list of protocols is longer in the display filters menu than in the capture filters menu. For display filters, the list includes all the protocols that your protocol interpreter suites can recognize.

Higher Level Addresses

In addition to the DLC source and destination addresses, a frame may contain packets that have their own, higher-level addresses. The interpreters display the higher level addresses and interpret them according to a name table that you can edit as you find convenient.

During display (although not during capture) you can set filters for specific higher-level addresses.

Address Level Filtering

An *address level* filter lets you accept frames only if they include an address at a particular protocol. When you set an address level filter, you check the various levels that you wish to include. A frame is accepted if one or more of the levels you checked is represented in its addressing (or in the addressing of an embedded packet).

For example, on an Ethernet network using NetBIOS, every frame has a DLC address. Some frames also have a NetBIOS address. In the address level filter, if you put a check mark beside DLC, the filter accepts every frame, since every frame has a DLC address. However, if you check NetBIOS but not DLC, the filter accepts only those frames that contain a NetBIOS address.

In the summary view, the analysis server shows each frame's address at the highest of the levels you checked.

Searching for Patterns in the Displayed Frames

You can search through the displayed frames for patterns, or combinations of patterns. Patterns are described in the same way as capture filters (page 1–8) or the trigger pattern that stops capture (page 1–12).

Searching for Text that Appears In Any of the Three Views

You can search through the capture buffer for a frame whose display contains a particular piece of text. Notice that the search is not limited to data actually present in the frame itself, but can also be for text that occurs in the summary or detail display that the analyzer generates when it describes the frame. This means that you could search for a symbolic address even though the name you look for is one you invented yourself and occurs nowhere in the transmission. Or you could look for a phrase used by the interpreter to describe the meaning of a code (for example, "file busy" or "access denied") even though the frame itself contains only a numeric value at a particular position.

Configuring Alternative Sets of Protocol Interpreters

Each Sniffer analysis server's software is compiled at the factory to match the type of network it will monitor (Ethernet, token ring, or WAN/synchronous). The software is also specific to the network and protocol stack that the analyzer will use to transport data and control information between the server and the console.

Each Sniffer analysis server contains copies of all current Network General protocol interpreter suites.

It is most unlikely that you would ever need all of the protocol interpreter suites at once. Moreover, the protocol interpreter suites are so large and so numerous that there isn't room for all of them at once in a single executable file. The factory therefore supplies *two* versions of the program called Sniffer Network Analyzer. You'll see two entries for Analyzer in the selection menu (Figure 1–6). (They'll all say Ethernet Analyzer, Token Ring Analyzer or WAN/Synchronous Analyzer, as appropriate.)

Between them, the two versions of the Sniffer analyzer contain all the protocol interpreter suites. Some interpreter suites are in one analyzer program, some in the other, and some in both. The division into two overlapping sets of protocol interpreters is initially provided at the factory. If you subsequently prefer other combinations of protocols, the Sniffer Configuration Utility (described in Chapter 7) makes it easy to generate them.

Identifying the Protocol Suites in a Particular Analyzer

When you highlight one of the "Analyzer" entries in the server's main selection menu, in the menu's lower panel you'll see a list of its protocol interpreter suites. That lets you distinguish between various Analyzers listed in the menu. (You can see the Analyzer protocol interpreter suites in the lower panel of Figure 1–6).

Suite	Name	Protocols
1301	IBM	SNA, SMB (including OS/2 LAN Manager), RPL, NetBIOS, IBMNM, SNAP, SDLC, LLC (802.2), MAC.
1302	Novell NetWare	NCP, SAP, NetBIOS, XNS, SPX, IPX, RIP, Echo, Error, AFRP, SNAP, LLC (802.2).
1303	XNS	SMB (including OS/2 LAN Manager), NBP, XNS, Courier, SPP, IDP, PEP, RIP, Echo, Error, SNAP, LLC (802.2).
1304	TCP/IP	SMB (including OS/2 LAN Manager), SNMP, Telnet, FTP, TFTP, SMTP, RUNIX, CMOT, DNS, TCP, UDP, IP, GGP, ARP, RARP, SNAP, TRLR, RIP, NetBIOS, ICMP, LLC (802.2), ASN.1.
1305	Sun	ND, NFS, YP, PMAP, Mount, RPC.
1306	ISO	X.400, FTAM (8571/4), VTP (9041), ACSE (8650/2), ISO Presentation (8823), ISO NetBIOS (MAP/TOP 3.0), ISO Session (8327), SMB (including OS/2 LAN Manager), TP (8073, class 0, 2, 4), CLNS (8473), ES-IS Routing (9542), SNAP, LLC (8802/2, type 1 and type 2), ISODE, ASN.1.
1307	DECnet	DAP, LAT, NICE, SMB, CTERM, FOUND, SCP, NSP, DRP, MOP, SNAP, LLC (802.2).
1308	Nestar	Nestar FS and IOB commands; SMB; XNS; SPP; IDP; FRP.
1309	Banyan VINES	StreetTalk, Mail, Route, Echo, Matchmaker, IPC, SMB, SPP, RTP, ARP, ICP, IP, FRP, SNAP, LLC (802.2).
1310	AppleTalk	AFP, TOPS, PAP, ASP, SoftTalk, ADSP, NBP, ATP, RTMP, ZIP, Echo, KSP, AARP, DDP, SNAP, LLC (802.2), LAP.
1311	X Windows	X Windows version 11, release 4.
1312	X.25	PAD (X.3, X.28, X.29), QLLC, SNDCP, X.25 layer 3, PPP, HDLC, SNAP, LLC (802.2).

Figure 1–5. Protocol Interpreter Suites.

The Analyzer's Menus and Controls

While you operate a Sniffer analyzer, you're always working from a SniffMaster console. There are menus to operate the console and menus to operate the individual servers. From the SniffMaster console menus, you activate the connection to one or more servers.

What You See When You Connect to a Server

When you display a server's screen, you've tuned in to what the server is doing at the moment. The server has a life of its own. It was running before you connected to it. It will go on running after you disconnect from it. It goes on running whether or not you choose to show its screen at your console. (You might think of this as coming into a room in which the server has previously been running with no one watching, but—now that you're in the room with it—you may choose to look at its screen.)

If the analysis server has been collecting data unattended, you see the display that was last requested, updated to show the current situation. If the analysis server has been started but given no specific instructions, you see its initial selection menu. If an earlier instruction exited from the server's menus and returned to DOS, you see the DOS prompt.

The Server's Selection Menu

When the analysis server is newly installed, or whenever it has just been reset or powered up, upon connection you'll see its *main selection menu*. From the selection menu, you can tell the server to run the Sniffer *monitor* or one of the Sniffer *analyzers*. You may also select the file transfer utility or to configure the server.

You can see an example in Figure 1–8. The selection menu lists the major choices open to you. (The items in the menu depend on the software modules and protocol interpreter suites installed in your particular Sniffer server.)

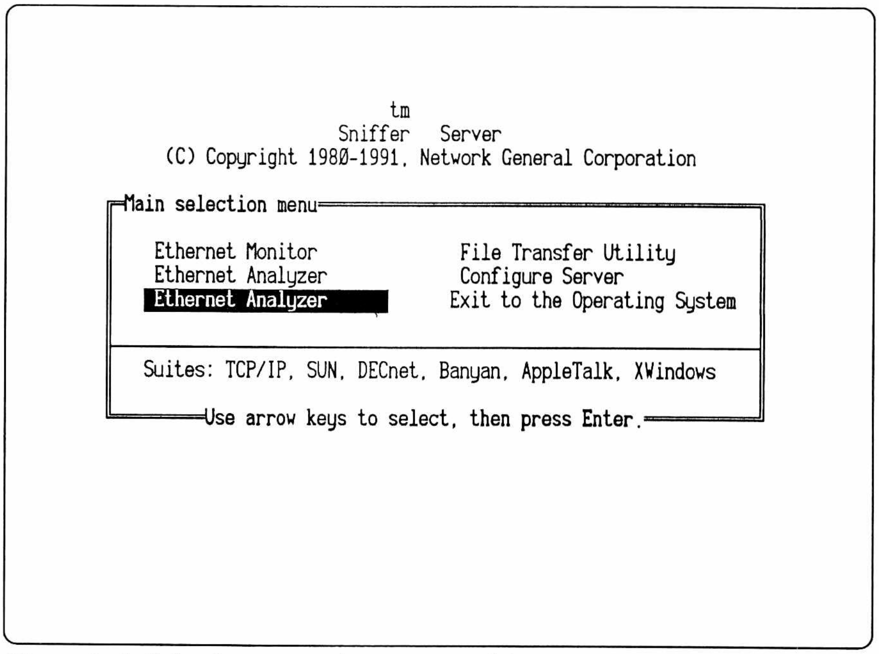

```
                         tm
                  Sniffer   Server
   (C) Copyright 198Ø-1991, Network General Corporation

 Main selection menu
    Ethernet Monitor                File Transfer Utility
    Ethernet Analyzer               Configure Server
    Ethernet Analyzer              Exit to the Operating System

   Suites: TCP/IP, SUN, DECnet, Banyan, AppleTalk, XWindows
          Use arrow keys to select, then press Enter.
```

Figure 1–6. Selection menu shows protocol interpreter suites for the highlighted analyzer

In each menu, one item is *highlighted.*[1] The lower portion of the panel contains a brief explanation of the highlighted entry. Pressing one of the cursor keys moves the highlight to another entry. To execute the entry that's highlighted, press Enter.

To start an analyzer

1. From the SniffMaster console, connect to a server. (For details regarding this step at the console, see *Distributed Sniffer System: Installation and Operations Manual*).

   Result: You see the server's current screen. When the server has been newly installed or has been reset, you'll see its main selection menu.

2. In the server's main selection menu, move the highlight to the analyzer you want and press Enter. (As you move the highlight to each analyzer, the list of its protocol interpreters appears in the menu's lower panel.)

Other Choices in the Server's Selection Menu

Some entries are always present in the main selection menu. For example, the menu always includes the option to exit to the operating system. The possible selections and the actions they perform are listed in Figure 1–7.

[1] We use "highlight" to mean the distinctive display of the selected item at the center of the screen. Depending on the type of display you're using, this may be in a contrasting color, or flashing or in reverse video.

Selection	Action
***xx* Monitor**	Starts the Sniffer monitor on the server, and passes control to the monitor's main menu.
***xx* Analyzer** (initially, two of them in the menu)	Starts the Sniffer analyzer on the server, and passes control to the analyzer's main menu. (The lower panel lists that analyzer's protocol suites.)
File Transfer Utility	Starts the server's end of the file transfer. When you invoke the corresponding utility at the console, you can transfer a file between them. For details, see *Distributed Sniffer System: Installation and Operations Manual.*
Configure Server	Brings up a menu in which you may elect: • To set server parameters. For details, see *Distributed Sniffer System: Installation and Operations Manual.* • To create (or delete) analyzers with particular combinations of protocol interpreter suites. For details, see Chapter 7.
Exit to the Operating System	Exits from the selection menu and returns to a subset of the DOS operating system. (The server is now inactive as an analyzer. If it was previously running the monitor and you did not explicitly shut down the monitor, the monitor continues to run in the background). To restart the menu, at the DOS prompt type **menu**.

Figure 1–7. Effect of items in the analysis server's main selection menu.

Connecting When the Analyzer or the Monitor is Already Running

When your console starts displaying the server's screen, you may find that the Sniffer analyzer is already running. (Or, on Ethernet or token ring, the Sniffer monitor may be already running.) If so, instead of the server's opening selection menu, you may immediately see a display from the analyzer (or the monitor, as the case may be).

If the server has been configured to accept more that one console, it's possible that, when you start your session, someone else is already connected from another console. Or some one else may connect to the same server later in your session.

When two consoles are connected to the same server, the server transmits displays to *both* consoles. It also accepts input from both consoles. The server doesn't assign priority to one console or another.

(When you configure a console to permit concurrent sessions from different consoles, it's a good idea to establish your own conventions about who's going to do this, and when.)

Assuming it's OK to proceed with a server that's already running, if you are about to start a new session, you probably want to start by returning to the monitor or analyzer's main menu (by pressing F5).

If you want to use the analyzer and you find that the the server is already running it, you can proceed with your work.

If you want to use the analyzer but you find that the server is at present running the monitor, you will have to stop the monitor software before you can start the analyzer.

To start an analyzer when the server is running the monitor

1. Press F5 to return to the monitor's main menu, and there select **Exit**.

   Result: That returns you to the server's Selection Menu.

2. From the selection menu, select **Monitor**.

3. In the Monitor's initial menu, select **Shutdown the Background Processes**.

   You'll see a warning message from the server, concluding with
   `Do you wish to shut down the monitor? [y/n]`
   Assent by typing `y`.

   When the monitor has shut down, you'll be back at the Selection Menu, but without the Monitor's background process running. Now you can start the analyzer.

4. Move the highlight to the analyzer you want and press Enter. (As you move the highlight to each analyzer, the list of its protocol interpreters appears in the menu's lower panel.)

Tree-Structure of the Analyzer's Menus

The Sniffer analysis server is entirely controlled from menus presented on the screen of the SniffMaster console. When you choose to run an analyzer, control passes immediately to the analyzer's main menu. An example (for an Ethernet analyzer) is shown in Figure 1–8.

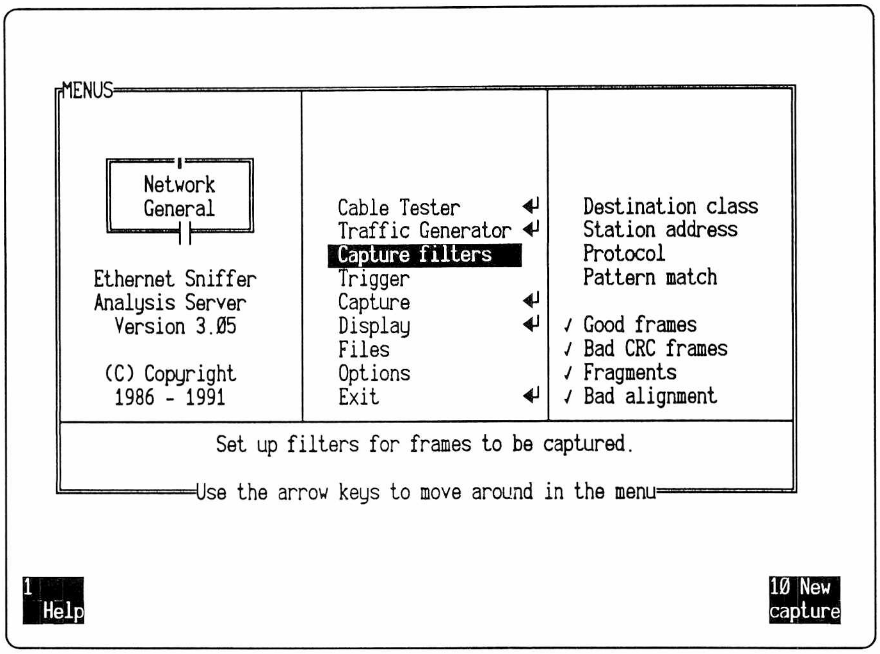

```
MENUS
  Network                Cable Tester        ↵   Destination class
  General                Traffic Generator   ↵   Station address
                         Capture filters         Protocol
Ethernet Sniffer         Trigger                 Pattern match
Analysis Server          Capture             ↵
 Version 3.05            Display             ↵   / Good frames
                         Files                   / Bad CRC frames
 (C) Copyright           Options                 / Fragments
  1986 - 1991            Exit                ↵   / Bad alignment

               Set up filters for frames to be captured.
            Use the arrow keys to move around in the menu

1                                                        10 New
 Help                                                    capture
```

Figure 1–8. Main menu of the Sniffer analyzer (in this case, for Ethernet).

All Sniffer analyzer menus have the same structure. While details vary according to the network and protocol interpreter suites installed, the organization is the same. The entire menu is a tree, with its root (the main menu) to the left and its branches and leaves to the right. Only a part of the menu is visible at a time. Using the cursor keys, you move a window over the menu. Since the window is fixed on the screen, the menu appears to move under the stationary window.

When the menu is active, the screen shows three panels side-by-side. You control the center panel. Within that center panel, the center row is highlighted. That is your location on the menu. When you first start a Sniffer analyzer, the center panel lists the alternatives available from the root of the tree. Some alternatives appear above the highlight, some below it. Initially, the highlight is on capture.

When you press the Cursor Up key, the item above Capture becomes highlighted. We speak of "moving the highlight up" to the next item. The highlight doesn't really move; instead, the entire center panel scrolls downward so that the (stationary) highlight is now on the row above. Alternatively, to jump directly up or down to a particular item, you may type the first letter of its name. The highlight jumps to the next item beginning with that letter.

The panel to the right shows choices in the submenu that go with the item highlighted in the center panel. As soon as you bring a different item to the center highlight, the entire right panel changes. The right panel always shows the submenu that goes with the item that's highlighted in the center (see Figure 1–9).

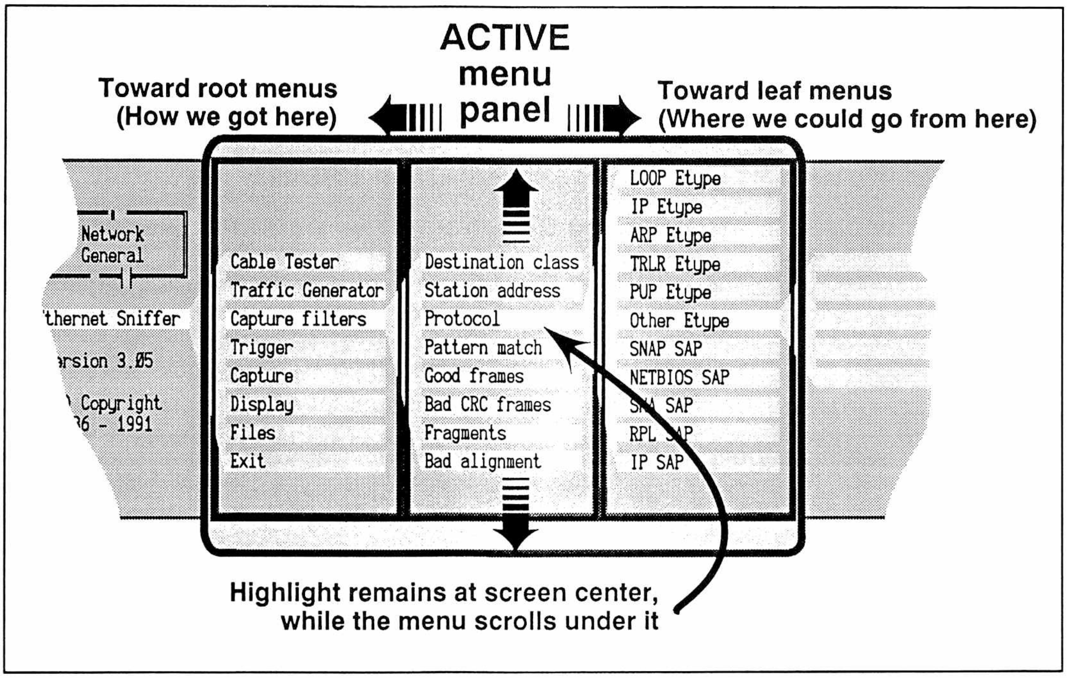

Figure 1–9. Scrolling over the tree-structured menu.

To select one of the options in the right panel, press the Cursor Right key. It's as if you move rightward from the center of the screen. The highlight doesn't really move. Instead, the entire menu moves leftward under the highlight panel. The panel that was to the right now moves to the center. The panel that was in the center now moves to the left. The panel that was at the left vanishes, as though it moved out of sight beyond the left edge of the screen.

The options associated with the item in the center appear in the panel to the right. As you move the highlight to the right, submenus for the next level seem to move in from beyond the screen's right edge. When the item in the center has no options, the right panel is blank.

The four cursor keys permit you to traverse the entire menu tree. Your current selection always appears at the center of the center panel, with other choices at the same level above and below it. The right panel shows the submenus (if there are any), and the left panel shows the menu nearer the root (or the Network General logo when you're at the root).

Menu Conventions

The menus contain two kinds of items:

- Options, and
- Actions.

First you set options; then you start an action involving them.

Options

An option specifies how an action will work when it starts at some future time. Examples of options:

- Choosing protocols to be included in the filters that decide whether to accept a frame.
- Setting a logarithmic scale for the traffic-density bar graph.
- Determining whether frames are displayed with the highest-level protocol only.

There are two kinds of options. A *checklist* is a list of items. Put a ✓ beside each that you want, an x beside those you don't want. You can check as many or as few as needed. A *radio control* is a list of mutually exclusive options (so called because, like a push-button radio, selecting one deselects the others). A radio control has a vertical bar beside it, and an arrowhead at the selected item, thus:

|▶ Selected item
| Deselected item

To change an option

Position the highlight on the option you want, and press Spacebar.

On a radio control, that moves the arrowhead to the highlighted item. On a checklist, pressing Spacebar changes ✓ to x (or x to ✓). Holding down the Alt key while pressing Spacebar reverses the setting of all items on the list.

Actions

An action that can be started by pressing Enter always has ↵ beside it.

When you press the key for an action, something starts to happen. For example, the action **Capture** causes the Sniffer analyzer to start capturing traffic according to the capture options you previously set. The action display causes the Sniffer analyzer to start displaying frames in the capture buffer according to the display options you previously set.

To start an action

Position the highlight on the action you want and press Enter.

CHAPTER TWO: MONITORING TRAFFIC AND CAPTURING FRAMES 2

Network General

Chapter 2. Monitoring Traffic and Capturing Frames

Chapter Overview

This chapter describes the process of capturing frames and the displays the analyzer generates as capture proceeds.

During capture, the Sniffer analyzer "listens" to all frames, filters them to select frames to record, and stores the selected frames in its capture buffer. As this process continues, the analyzer updates graphs or tables that summarize the accepted frames.

Before you start capture, you need to set a number of options that govern the frames that will be accepted and the conditions under which capture will halt. The chapter reviews these in the order that they appear in the Sniffer analyzer's menus.

When capture is complete, you can then analyze, interpret, and display the captured frames, as described in Chapter 5, "Displaying and Interpreting Frames."

Monitoring During Capture vs. Full-Time Monitoring

While it is capturing, the analyzer counts the arriving frames and generates some displays. This activity amounts to some rather restricted monitoring. More complete facilities for monitoring are provided in the separate Sniffer monitor application. It devotes full attention to monitoring and doesn't attempt to save frames. Thus, to get maximum monitoring but no capture, you should run the Sniffer Monitor. See the separate manuals *Distributed Sniffer System: Ethernet Monitor Operations Manual* and *Distributed Sniffer System: Token Ring Monitor Operations Manual.*

The Analyzer's Process of Capture

Each accepted frame is stored in the capture buffer. You can store an entire frame or truncate it after a certain length.

If you keep on capturing after the buffer fills, earlier frames are discarded to make room for later ones.

When capture is terminated, you can display and analyze the frames in the buffer. That's when filtering and decoding of higher-level protocols takes place. You can also store the contents of the capture buffer in a disk file, and subsequently analyze it, or play it back as though it were being captured again.

Displays During Capture

Skyline A graph showing traffic density over time; or

Counters A table of counters, perhaps by logical call and type (on a synchronous WAN) or by source or source and destination (on a LAN).

List of Preparations for Capture

Before you start capture, you specify the frames you want the analyzer to retain in its capture buffer and the displays you want to see while frames are being captured.

Many options affect the way the analyzer captures frames. Generally speaking, you must set these conditions before you start capturing. (However, when you've loaded a setup specifying everything you want, or you're happy with the defaults, you can start capture immediately, without going through these preliminaries.)

Figure 2–1 is a list the topics involved in preparation for capture. Each of these is described in more detail in the sections that follow.

Preparations for Capture

Check options	Several general options affect the analyzer's operation. They're set in a menu entitled "Options."
Set capture filters	Filters tell the analyzer which frames to retain and which to discard. A capture filter can be based on the presence or absence of the following: • Unknown address (LAN) • Destination class (LAN) • Low-level protocol (WAN) • Direction (DTE/DCE) (WAN) • Pattern (both LAN and WAN) • Defects (both WAN and LAN.)
Set the trigger	When it sees a frame containing the pattern, the analyzer stops capturing. Preparation has two parts: • Set the pattern to look for • Set the stopping point (how much to capture) after it's found.
Edit the name table	During capture, the analyzer labels sources and destinations using names from the name table. The unknown station filter considers a station "unknown" if its address has no name in the table.
Set the source	The source of capture may be either "live" from the network or from a file of previously recorded frames.
Set the frame size	The default is to record the entire frame. To save space, you can have the analyzer truncate each frame.
Set display type	• Counters • "Skyline" histograms.
Set counter units	• Frames seen • Kilobytes seen • Network utilization.
Set graph scale	The traffic density bar graph's horizontal scale may be: • *Logarithmic* (default) • *Linear.*
Set skyline interval	Update the skyline at the end of each: • Second • Minute • Hour.
Set "Highspeed"	On Ethernet, you can suppress capture displays to reduce the risk of losing frames during a highspeed burst.
Start capture!	• Press F10 (labeled New Capture) • Highlight **Capture** and press Enter.

Figure 2–1. Steps in preparation for capture.

Choices, Setups, and Defaults

For each step, you may indicate what you want, or (by inaction) accept what is already established. "Established" may mean either

- The Sniffer analyzer's default setup, or
- The settings you have restored by loading a previously-saved setup file. Loading a setup file restores all the settings the way they were at the time the setup file was saved. Procedures to save and load a setup file are described in Chapter 5, starting at page 5–71.

Setting the General Options for Each Analyzer

The Options menu is at the bottom of the main menu's first panel. You probably won't need to change these settings often. They reflect general characteristics of the network or the analyzer.

Token Ring Option to Remove Itself Automatically

What should the analyzer do if it receives no signal when it inserts into the token ring? You may select between two choices: remove the server from the ring immediately or remain. The default is to remove, indicated in the options menu by the setting:

```
√ No signal: remove
```

Leaving immediately is a matter of prudence. It protects the ring if you should inadvertently connect a server that has been configured for one speed to a ring that is operating at a different transmission speed. Connecting a token ring station at the wrong speed may severely disrupt the ring.

If you're certain that the server's speed matches the speed of the network, you can remove this protection. Operating without automatic withdrawal has two advantages:

- When the ring has been broken, you can connect to a fragment and await signals as you make other changes to the ring.
- On a functioning ring, you can disconnect temporarily and reconnect without having to reset the Sniffer server.

By moving the highlight to **No signal: remove** and pressing Spacebar, you reverse the √ (active) to X (inactive).

LAN Option to Interpret or Ignore the RI Bit in a Source Address

This option affects the way the analyzer treats DLC addresses on networks that use six-byte addressing (including token ring and

Ethernet). Some networks reserve one bit in the source address to indicate that the frame includes a field called *source routing information* (RI). On networks that recognize the RI bit, an address really consists of only 47 bits. The 48th bit is used in the destination address to indicate "broadcast" and in the source address to indicate "RI present." The broadcast bit is the one that is physically transmitted first.[1]

Effects of enabling "Interpret RI"

- The analyzer treats addresses as 47 bits
- The analyzer calls the RI interpreter to interpret the RI field of frames that contain this bit in the source address.

A few networks treat all 48 bits as part of the address. On such a network, the 48th bit is simply part of the address, and does not mean that there is an RI field. When analyzing frames from such a network, it's important to disable the option to **Interpret RI**. (As shipped, the analyzer's default has **Interpret RI** enabled.) If you don't turn off the interpretation of the RI bit, the analyzer may try to interpret part of the frame's data field as an RI field.

Effects of disabling "Interpret RI"

- The analyzer treats addresses as 48 bits
- The analyzer does *not* call the RI interpreter and assumes that there is *never* an RI field.
- The Sniffer analyzer's "destination class" filter continues to recognize the "broadcast" bit in a destination address. Thus the filter treats a station whose address includes a 1 in that position as if it were a broadcast.

To treat the high-order bit of the source address as part of the address, and not an RI indicator

1. From the main menu, move to **Options**.
2. Move to **Interpret RI**. Press Spacebar to change ✓ (active, the default) to X (inactive).

RI Fields and "Data Relative" vs. "Frame Relative"

The IEEE 802.2 standard does not mention RI. IBM introduced source routing on token ring networks, and that remains the context in which it is most frequently found. In principle, source routing information

[1] Because there are different rules for converting bits-on-the -wire to bits-in-memory, the broadcast bit is the high-order bit of a token-ring address, but in an Ethernet address it's the low-order bit of the first byte.

can be used not just on token ring but on any LAN that uses six-byte station addresses, including Ethernet, StarLAN, or PC Network.

RI is a variable-length field inserted after the DLC destination and source and before the frame's usual data. When a frame contains an RI field, the remaining data field is displaced by the total length of the RI field. To allow for the varying size of the RI field, the Sniffer analyzer permits you to describe a pattern by either its frame relative or its data relative offset; see page 2–25.

The RI field contains an identifier for each of the routers or gateways that forwarded the frame. As it retransmits a frame, each router appends its own two-byte identifier[1] to the RI field. By this means, the ultimate recipient has a record of all the intermediate stations that forwarded the frame.

To request that each intermediary insert a record of itself, the originating station turns on the RI bit. That instructs each station that receives the frame to interpret the first two data bytes as the RI header, perhaps followed by a list of gateway identifiers.

The originating station sets up the RI header. Five bits in the RI header record the total length of the RI field (including the header). Initially, when no gateway has yet forwarded the frame, the total length of the RI field is 2 bytes. Each gateway that forwards the frame appends its own two-byte identifier to the list of gateways and increases the RI length by 2.

WAN/Synchronous Options

WAN Packet type

You have to tell the Sniffer analyzer what lower-level protocols the synchronous link uses. The lowest level may be either SDLC (IBM's Synchronous Data Link Control protocol), or HDLC (High-level Data Link Control, an ISO standard protocol descended from and closely related to SDLC). Neither of these affects what higher level protocols are embedded in their frames. The most widely used are SNA over SDLC (in IBM installations), and X.25 over HDLC (widespread in Europe and increasingly in the US). The analyzer offers two choices:

HDLC/X.25	ISO X.25 over HDLC (the default).
SDLC/SNA	IBM's SNA (Systems Network Architecture) protocol over SDLC.

[1] The identifier is composed of a 12-bit ring number and a 4-bit bridge number. These are arbitrary numbers assigned by the network administrators. The RI identifier is unrelated to the bridge's station address.

WAN Packet Numbering

Packets (frames) may have sequence numbers generated in either of two ways that are not readily distinguishable by inspection. One uses 3 bits, the other 7. Three-bit (that is, modulo 8) numbering is widely used within the US and Europe, but seven-bit (modulo 128) encoding is often used in Japan and in international satellite links. The menu choices are:

▸ Modulo 8 (the default).
Modulo 128

WAN Data signalling

The two most common encoding methods for SDLC / HDLC are called NRZ and NRZI. Before you start capture, in order to decode the transmitted data correctly, you should tell the analyzer which one the network uses. The choices are:

▸ `NRZ` Non-return to zero (the default.)
`NRZI` Non-return to zero inverted.

Names in Filters and Displays During Capture

In any of the tabular displays during capture, if the analyzer's name table includes a symbolic name for a DLC address or for the source or destination of a logical call, it displays the name rather than the address.

Similarly, any filter that refers to a station's name also uses the symbolic version if the name table provides one.

To include symbolic names in displays during capture

1. Before setting filters or capturing, edit the name table to include names for the address you expect. See page 5–62 for the procedure to edit the name table.

   Alternatively...

1. Capture for a while without an updated name table. During capture, the analyzer will display numeric addresses for stations it doesn't recognize.
2. Display the captured frames (see page 5–5). When you start display, the analyzer scans the capture buffer for addresses and lists them in the working copy of the name table.
3. During the display, if possible, execute **Look for names** (page 5–66) or **Resolve names** (page 5–66).

4. Before closing the session or discarding the captured frames, execute **Edit names** (page 5–62). Provide names for all addresses.
5. Before exiting, execute **Save names** (page 5–61).
6. Start a new capture session. The analyzer will then display the names thus added to the name table.

Setting the Capture Filters

From the main menu, move upward until the highlight is on **Capture filters**. At the right you'll see the sub-menu of capture-filter options (visible in Figure 2–2). Move the highlight to the right and then up or down as necessary to display or edit the options.

When you specify a capture filter, you are setting the current capture filter, stored in main memory. Doing so does not update any stored file containing your Sniffer analyzer's setup, so the revision is temporary. When you have completed your revisions, you can, if you wish, save a record of the setup. When you subsequently load the setup file, your entire setup will be restored, including whatever filters you had set. To save your current setup, move to **Files** in the main menu, then to **Save**, and finally to **Setup**. (See page 5–70 ff. For more information about saving and loading setup files.)

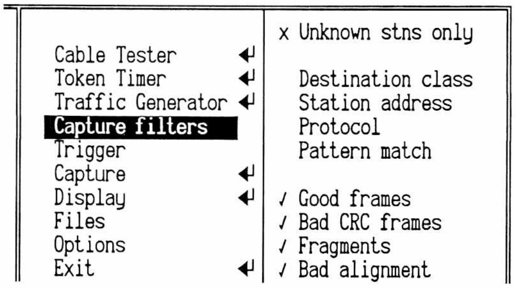

Filters for address, unknown, station, or destination class exist only on a LAN.

Filters for fragments or alignment exist only on Ethernet.

Filters for Bad CRC exist only on Ethernet or WAN/Synchronous.

Figure 2–2. Choices in the "Capture Filters" menu. Some items may be omitted depending on the network.

Five Kinds of Capture Filters

Figure 2–3 lists and describes the kinds of capture filters. A frame is accepted if it is not rejected by any of the applicable filters. (Put another way, a frame is accepted only if it passes all filters.)

Network	Filter	Test	Default
LAN only	Unknown station	Should the analyzer accept a frame only when it contains a source or destination that has not been assigned a symbolic equivalent in the current name table?	No (i.e. no restriction)
	Destination class	Does the frame have a specific destination or is it broadcast?	Accept both
	Station address	Does the frame match any of a list of source-and-destination pairs?	None set (accept any)
	Protocol	Does the frame contain any of a list of low-level protocols?	None set (accept any)
	Frame defects	Should the analyzer accept a frame that is defective, for example, having • Alignment error, or	Accept
Both LAN and WAN		• Bad CRC.	
	Pattern match	Does the frame match a set of patterns you have specified?	None set (accept any)
WAN only	Direction	Should the analyzer accept a frame coming • From DTE • From DCE	Accept both
	Flow control	Should the analyzer accept • RR frames (receiver ready) • RNR frames (receiver not ready)	Accept both

Figure 2–3. Types of capture filters.

Time Required for Capture Filters

Capture filters take a certain amount of processing time. In general, the more complex the capture filter, the greater the time. On a network whose load is light or moderate, the processing time is not noticeable. But on a heavily loaded network, it limits the speed at which the analyzer can accept frames, unless the filter significantly reduces the number of frames accepted. If frames arrive faster than the analyzer can process them, some may be lost. The number lost is recorded in the *lost frames* counter. If this number begins to rise, it may help to adopt a simpler filter.

Unknown Address Filter (LAN)

The *unknown address* filter selects frames whose source or destination is unknown to the analyzer. The cause of such traffic could be illegal hackers, bad data frames, faulty software in the application or the network, or network interface cards that have been replaced without notifying the system manager.

When you check **Unknown stns only**, the analyzer accepts a frame when it contains a DLC address (source or destination) that is "unknown." The analyzer considers an address to be "unknown" when

- Its address isn't in the working name table, or
- Its address is in the table but without a name.

To restrict capture to frames from unknown stations

1. In the Capture filters menu, highlight **Unknown stns** only, and press Spacebar to toggle between √ (active) and X (inactive).
2. In the **Display** menu, highlight **Manage names**. Then select **Edit names**. Edit the table so that
   — It includes the addresses of known stations;

   — Each known address has a symbolic name.

When you first display frames from a new capture, the analyzer scans the capture buffer for DLC addresses that are not in the name table, and automatically inserts them in the working copy of the table. However, it does not update the permanent name table (file STARTUP.xxD) unless you execute **Save names**. Even then, it saves only those addresses that have names. Thus, to assure that observed addresses are retained, you must both name each address and execute **Save names**.

Since the name table is mostly concerned with display, it is described in Chapter 5. For instructions on editing the names in the name table, see page 5–62.

Destination Class Filter (LAN)

On a LAN, a frame may either be directed to a specific destination, or to a generic address that several—perhaps all—stations accept. There is no equivalent type of transmission over a wide area network, and hence no destination class filter for a WAN/synchronous link.

In the name table, you can assign a name to a generic address, for example "Error Monitor" or "LAN Manager." Of course, a generic address will turn up only as a destination, never as the source of a frame. In each network, the prescribed form of a generic address is such that it is always different from any possible individual address.

Networks permit different types of generic addresses and identify them differently. The possibilities are summarized in Figure 2–4.

Originating Network	Type	Description	Characteristic Address
Ethernet, StarLAN, PC Network	Multicast address (including Broadcast)	A multicast address is a collective name for several stations. It may be a role played by one or more stations, or by all stations (for example, "Broadcast").	A DLC multicast address has a 1 in the low-order bit of the first byte, so that in hexadecimal its second character is odd (that is, 1, 3, 5, 7, 9, B, D or F). No individual station has an address with that bit on.
Token ring	Functional address (including Broadcast)	A functional address is a collective name for a role played by one or more stations, by no station (for example "Error monitor" when no station is monitoring errors), or by all stations (for example, "Broadcast").	A DLC functional address has a 1 in the high-order bit, so that in hexadecimal it appears as a number whose first digit is 8 or more. No individual station has an address with that bit on.
WAN/ Synchronous	None		

Figure 2–4. Types of generic addresses, by network of origin.

To select or exclude frames by destination class

1. Note the meaning of broadcast address in the network you are monitoring.

2. Move the highlight to **Capture filters** and then to **Destination class**. Press Spacebar to select or deselect
   √ Broadcast
   √ Specific.

Station Address Filters

During capture, the Sniffer analyzer's address filters consider only the lower levels of addressing. On a WAN/synchronous link, there are only two low-level addresses: the two ends of the link. In that context, filtering by DLC address doesn't make sense, and there is no provision for address filters during capture from a WAN.

On a LAN, you can set filters for a frame's DLC destination. However, recall that the DLC destination may describe only the current leg of a much longer journey. Higher-level protocols embedded in the frame's data field will have their own addresses; they may cause the recipient of the current frame to repack the data and retransmit it with a new address.

During capture from any of the LANs, you may specify up to four source-and-destination pairs for the capture filter. You may also specify what you want done with *others*—that is, frames that don't match the specifications you provide.

Naming the "Matches" Used in Setting Up a Filter

Each source-and-destination pair that you describe is called a *match.* To make the various matches easier to describe, you may assign them names. (The names don't do anything except help you remember what the various matches are for.) Initially, the matches are called match 1 through match 4. To assign a name to one of them, move the highlight to **Match 1** (or whatever) and press Enter. The Sniffer analyzer opens a dialog box in which you can write a name (Figure 2–5). Thereafter, you'll see the name you enter in place of the name Match 1.

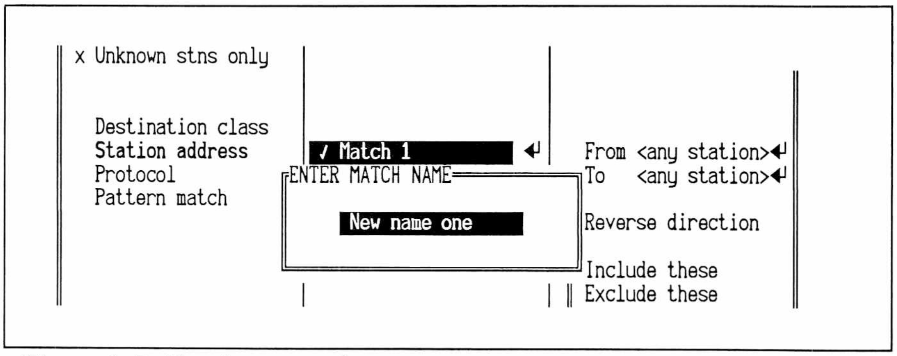

```
x Unknown stns only

Destination class
Station address      ✓ Match 1          ↵   From <any station>↵
Protocol            ENTER MATCH NAME          To   <any station>↵
Pattern match
                       New name one           Reverse direction

                                              Include these
                                              Exclude these
```

Figure 2–5. Naming a match.

For each match, you specify:

- Source (select a station name from the name table);
- Destination(select a station name from the name table);
- Whether to include frames traveling between the same pair but in the reverse direction;
- Whether the filter should include or exclude the frames thus identified.

To set station address filters for capture

1. Move the highlight to **Capture filters**, then to **Station address**, then rightward to **Match 1**.

2. Optional: To provide a name for this match, press Enter. A dialog opens to receive the name. Supply a name, and press Enter.

3. To specify a source address, move the highlight to **From** and press Enter. The analyzer opens a dialog box headed SELECT STATION. It contains the list of DLC addresses and their current names (Figure 2–6). Move the highlight vertically until it shows the station you want, and press Enter.

   Result: at the position labeled **From**, the server inserts the address, using its name if it has one, or the numeric address otherwise.

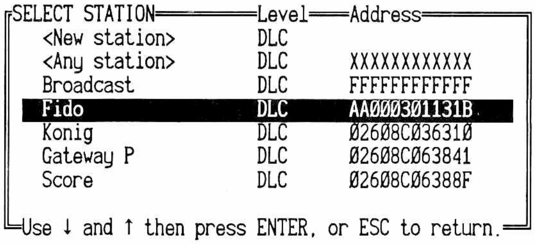

```
SELECT STATION════════════Level════Address══════════
  <New station>           DLC
  <Any station>           DLC      XXXXXXXXXXXX
  Broadcast               DLC      FFFFFFFFFFFF
  Fido                    DLC      AAØØØ3Ø1131B
  Konig                   DLC      Ø26Ø8CØ3631Ø
  Gateway P               DLC      Ø26Ø8CØ63841
  Score                   DLC      Ø26Ø8CØ6388F

Use ↓ and ↑ then press ENTER, or ESC to return.
```

Length of the hex station address depends on the network

Figure 2–6. Menu to select a station for a station address filter.

4. If the address you want isn't in the table, move the highlight to **<New station>** and press Enter. The analyzer opens a dialog box to receive your description of the new station. Enter a new DLC address and then a symbolic name for it (Figure 2–7). Press Enter when done, and return to Step 3.

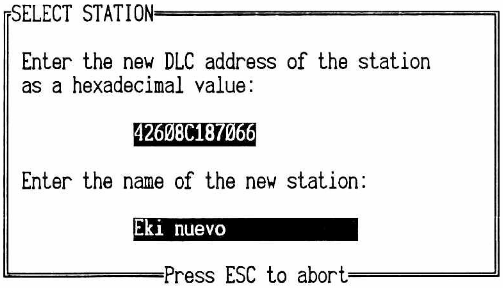

```
SELECT STATION═══════════════════════════════════

 Enter the new DLC address of the station
 as a hexadecimal value:

            426Ø8C187Ø66

 Enter the name of the new station:

            Eki nuevo

═══════════════Press ESC to abort════════════════
```

The dialog box requires an address of the appropriate length.

Figure 2–7. Dialog box for inserting a new station address and name.

5. To specify a destination address, move the highlight to **To** and press Enter. Supply an address in the same way as Step 3.

6. To indicate whether this match also applies to traffic in the reverse direction, move the highlight to **Reverse direction**. Press Spacebar to toggle between √ (include) and x (exclude).

7. To indicate whether frames identified by this match should be included or excluded, move the highlight to the radio control **Include these** or **Exclude these**, and press Spacebar at the one you want.
8. Repeat steps 1 through 8 for up to four matches.
9. To specify what the analyzer should do with frames not covered by your matches, move the highlight to **Others** and then to **Include** or **Exclude**, and there press Spacebar.

Combined Effect of Several Matches in an Address Filter

The Sniffer analyzer evaluates the matches in sequence from Match 1 (or whatever you have renamed it) to Match 4. As soon as a match succeeds, the server ceases to test that frame for further matches. The frame's fate is determined by the way you set the switch **include/exclude** for the match that succeeded.

When none of the four matches succeeds, the frame's fate is determined by the way you set the switch **include/exclude others**.

Using Fewer than Four Matches

It isn't necessary to set all four matches; indeed, you may not want to set any at all. In that case, don't set the matches you don't want. The analyzer disregards a match that contains only the default settings **Include** with addresses **From <any station>** and **To <any station>**. When all four matches are **Include From <any station> To <any station>**, the Sniffer analyzer does not use any address filtering during capture.

You can temporarily disable a match that you've defined by putting an X in front of its match name. The analyzer checks only the matches marked ✓. As usual, you toggle between ✓ and X by highlighting the match's name and pressing Spacebar.

Sample Address Filters

Suppose you're interested in traffic between George or Anita and Server-1, and also any traffic going to Gateway-A (but not the reverse). You could specify three matches as shown in Figure 2–8.

	Name for Match	Source Address	Destination Address	Reverse Direction	Include/ Exclude	
Match 1	G's files	George	Server-1	yes	include	
Match 2	A's files	Anita	Server-1	yes	include	
Match 3	Outgoing	<any>	Gateway-A	no	include	
Match 4	[Match 4]	[<any>]	[<any>]	[yes]	[include]	Defaults
Others					exclude	Default

Figure 2–8. Example of a filter-match on three address pairs.

Taking Advantage of the Order in Which Matches are Checked

Suppose you're interested in all traffic to and from Server-1 except for the voluminous traffic between Server-1 and Gateway-A, which, if included, would fill up the capture buffer with frames that don't at the moment concern you. Since the address filters are evaluated in sequence, you can readily achieve the desired effect by first discarding frames between Server-1 and Gateway-A and then accepting frames between Server-1 and any destination. The set of matches is summarized in Figure 2–9.

	Name for Match	Source Address	Destination Address	Reverse direction	Include/ exclude	
Match 1	S1 to A	Server-1	Gateway-A	yes	exclude	
Match 2	S1 to others	Server-1	<any>	yes	include	
Match 3	[Match 3]	<any>	[<any>]	]yes]	[include]	Defaults
Match 4	[Match 4]	[<any>]	[<any>]	[yes]	[include]	Defaults
Others					exclude	Default

Figure 2–9. Address filter to capture all traffic from a station, but with a major exception

Protocol Filter During Capture

While capturing from a synchronous link, there is no option to filter by protocol, so this section applies only to LANs.[1]

On a LAN, each DLC frame includes a field that indicates what it contains. Depending on the network, this low-level classification is called a *SAP* ("service access point" in the terminology of IEEE 802.2), or an *Ethertype* (on Ethernet and StarLAN). This is the lowest level at which protocols are identified, and the only level that can be used for capture filters. (During display, the analyzer can devote more time to processing filters, so display filters can include tests for higher-level protocols or addresses.)

When you move the highlight in the **Capture filters** menu to **Protocol**, the panel to the right shows you a list of protocols. Beside each name, there's a √ (check mark) if the capture filter accepts it and an X if it doesn't. The filter accepts a frame if it contains *any* of the protocols you have marked with a check.

The specific protocols in the list that accompanies the capture filters menu depend on the network for which your Sniffer analyzer is prepared and the protocol interpreter suites in the analyzer you are using.[2] Some typical protocol lists are shown in Figure 2–10.

[1] In principle, you could filter for SNA or X.25, but since these would not be intermixed on the same link at the same time, the filter would do nothing useful.

[2] When you highlight the word **protocol**, the lower panel lists the installed protocol interpreter suites.

Token ring with IBM, XNS/MSNET, ISO, X.25	Ethernet with TCP/IP, Sun, DECnet, Banyan, AppleTalk, X-Windows	Ethernet with IBM, Novell, XNS/MSNET, ISO, X.25
√ MAC frames √ SNAP SAP √ BPDU SAP √ NetBIOS (IBM) SAP √ SNA SAP √ RPL SAP √ U-B SAP √ IBMNM SAP √ NetWare SAP √ ISO CLNP SAP √ XNS SAP √ X.25 SAP √ Other SAP	√ LOOP Etype √ Com Netmap Etype √ IP Etype √ ARP Etype √ TRLR Etype √ PUP Etype √ PUP ARP Etype √ SNMP Etype √ MOP Etype √ DRP Etype √ LAT Etype √ IP (VINES) Etype √ LOOP (VINES) Etype √ Echo (VINES) Etype √ ARP (Atalk) Etype √ LAP (Atalk) Etype √ Other Etype √ SNAP SAP √ BPDU SAP √ LLC (VINES) SAP √ Other SAP	√ LOOP Etype √ 3Com Netmap Etype √ IBMRT Etype √ NetWare Etype √ XNS Etype √ 3Com NBP Etype √ PUP ARP Etype √ Other Etype √ SNAP SAP √ BPDU SAP √ NetBIOS (IBM) SAP √ SNA SAP √ RPL SAP √ IBMNM SAP √ ISO/NetWare SAP √ NetWare SAP √ X.25 SAP √ Other SAP

Figure 2–10. DLC protocols typically available for capture filters

To choose the protocols the analyzer will accept during capture

1. In the **Capture filters** menu, move to **Protocol**, and then rightward to the panel containing the list of protocols.
2. Move the highlight vertically to the protocol you want to select. Press Spacebar to toggle between √ (selected) and X (not selected).
3. To reverse the setting for all protocols, press Alt-Spacebar.
4. The last entry on the list of protocols is "other"; here indicate what should be done with a protocol that isn't any of those the analyzer recognizes.

For most networks, the analyzer's default setting is to accept every protocol.

Pattern Matching

A pattern is a particular sequence of bits within a frame. In a simple pattern, the bits occur at just one location. In a complex pattern, a set of up to eight simple patterns is linked by AND or OR; the effect of

NOT is achieved by a setting labeled **Match/Don't match**. The resulting set of patterns becomes one of the filters applied during capture.

Four Contexts for Pattern Matching

There are four different contexts in which the Sniffer analyzer lets you specify a pattern to look for. As you will see, the pattern can be quite complex, involving logical combinations of up to eight component patterns. The four contexts are independent; a pattern established for one of them has no effect on patterns established for any of the other three. However, the analyzer uses exactly the same mechanism for specifying a pattern in any of the four contexts. Therefore the discussion that follows, telling you how to set up a pattern for the capture filter, applies equally to setting a trigger pattern, setting a display pattern, or setting a search pattern. The discussions of pattern matching in three other parts of the manual therefore refer you back to here. The four contexts for pattern matching are:

Capture filter During capture, the analyzer accepts frames that contain the set of patterns specified in the capture filter.

Trigger Capture can be halted when the server finds that an accepted frame matches the set of patterns specified as the "trigger" pattern.

Display filter During display of frames in the capture buffer, you can restrict display to frames that match the set of patterns in the display filter.

Search During display, you can have the analyzer search forward through the capture buffer for the next frame that matches the set of search patterns.

Setting patterns in the capture filter offers a way of filtering for higher-level protocols. Recall that, during capture, the analyzer doesn't have time to invoke its interpreters for high-level protocols, and so can't directly filter on high-level protocols or high-level address. However, it does have time to execute fairly complex pattern matching. By a little experimentation with captured frames, you can often set up a pattern that in effect responds to a high-level embedded protocol.

Establishing a Complex Pattern by Combining Four Matches

The panel to the right of **Pattern match** shows the logical rules for combining the four sub-patterns into a set. To help recognize the four component patterns, you can assign each a name. Your name then replaces the default names Match 1, Match 2, etc. Your substitute names don't affect the processing; they only serve to remind you what each component pattern is for. You assign a name to a match in the same way that you name a station-address match (Figure 2–5, page 2–14). That is, move the highlight to Match 1 (or whatever) and press Enter. The analyzer opens a dialog box to receive your substitute name.

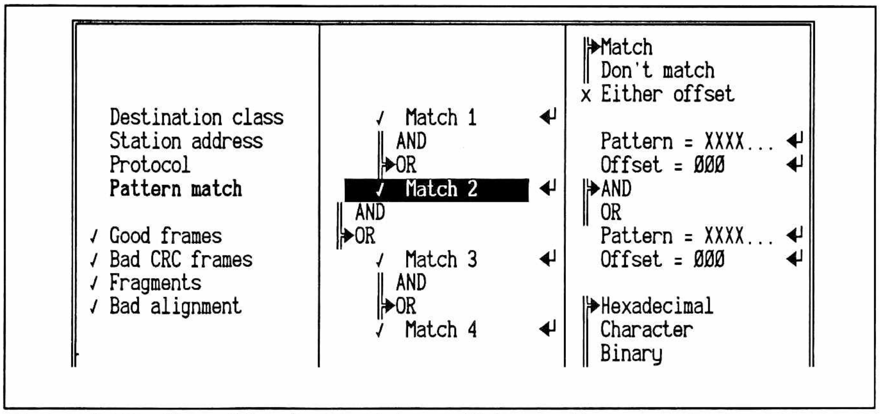

```
                                                          ▶Match
                                                           Don't match
                                                         x Either offset
   Destination class          ✓  Match 1         ↵
   Station address             AND                         Pattern = XXXX...  ↵
   Protocol                   ▶OR                          Offset = 000       ↵
   Pattern match              ✓  Match 2         ↵        ▶AND
                         AND                               OR
                        ▶OR                                Pattern = XXXX...  ↵
 ✓ Good frames                ✓  Match 3         ↵         Offset = 000       ↵
 ✓ Bad CRC frames              AND
 ✓ Fragments                  ▶OR                         ▶Hexadecimal
 ✓ Bad alignment              ✓  Match 4         ↵         Character
                                                           Binary
```

Figure 2–11. One of the four matches within a pattern.

Logical Combinations of the Four Matches

The four matches are combined by logical AND or OR (visible in the center panel of Figure 2–11). Initially, these are all set to OR.

The matches are grouped into two sets of two (Figure 2–11). The relation between Match 1 and Match 2 is evaluated first; then, the relation between Match 3 and Match 4; finally, the overall relation between the outcomes of those pairs. Algebraically (with the symbol ⊗ standing for whichever relation you set, either AND or OR), it's

$$(\textit{Match 1} \otimes \textit{Match 2}) \otimes (\textit{Match 3} \otimes \textit{Match 4})$$

The menu conveys that by its pattern of indents, like a tree with its root at the left and its leaves at the right (Figure 2–12).

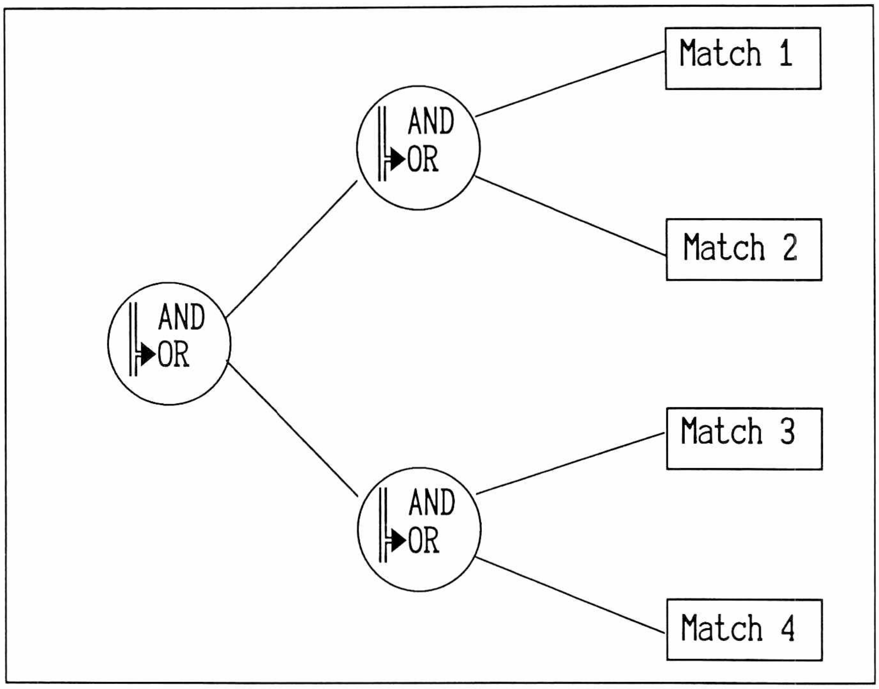

Figure 2–12. Combining four matches by AND and OR.

Pair of Patterns within a Match

Each individual match is in turn composed of a pair of patterns, also linked by AND or OR. (Those for Match 2 are visible in the right panel of Figure 2–11). When you highlight a match name, the panel to the right shows its pair of component patterns (Figure 2–11). Initially, all the individual patterns contain X (for anything), and all the offsets are 00.

A pattern whose characters are all X has no effect. When you don't need all eight of the individual patterns, you don't have to set the ones you don't use. Similarly, when you don't need all four of the matches, you don't have to set the matches you don't use. When you use some matches but not others, it doesn't matter which ones you use and which you leave at their defaults.

Temporarily Deactivating a Match

Beside each match, the symbol ✓ indicates that it is *active* while X indicates that it is *inactive*. Changing ✓ to X lets you deactivate a match temporarily without having to erase its specifications. By default, they're all marked "active."

Highlighting the match name and pressing Spacebar also serves to reactivate a match that you've deactivated.

A Match on a Pair of Patterns

A single match can involve a pair of patterns. (You may prefer to think of this as a pattern in two parts). In the menu (Figure 2–11), the two patterns appear one above the other, each with its offset. You aren't required to fill in both parts; any part you leave unspecified has no effect.

If you specify two patterns, you must also state the relationship between them: AND or OR. The default is AND. That is, to satisfy the match, the frame must contain *both* patterns.

For each pattern, you specify an *offset* (its location in the frame). Think of them as pattern A at offset *a*, and pattern B at offset *b*. Then the default match is:

Pattern A at offset *a* AND Pattern B at offset *b*

Since these two patterns are in the same frame, a frame that meets this condition looks like Figure 2–13.

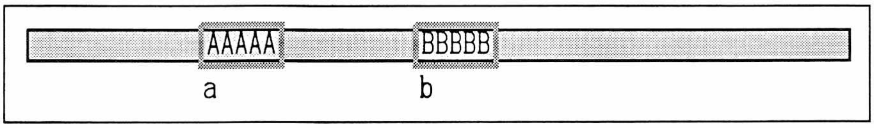

Figure 2–13. Frame containing both A at a *AND B at* b.

When the relationship is OR rather than AND, either of the frames shown in the top part of Figure 2–14 is acceptable. (And of course a frame that has both is still acceptable.)

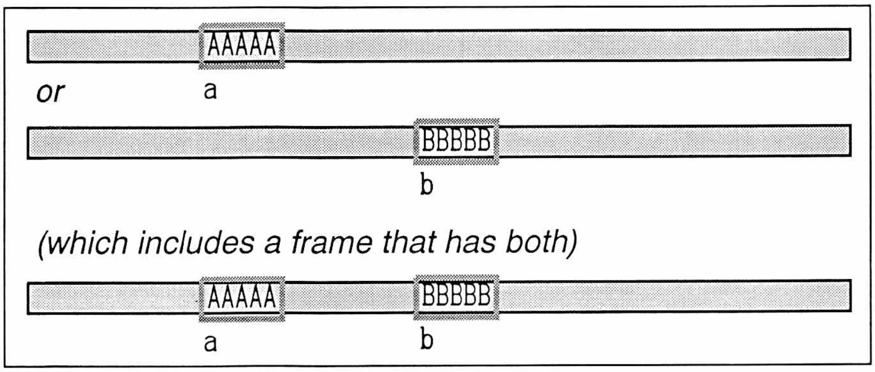

Figure 2–14. Frames containing either A at offset a OR B at offset b.

Either Offset

When you specify a pair of patterns, you also have the option to select **either offset**. This usually makes sense only when the two patterns are related by AND.

For example, an exchange between client and server over TCP might involve the "well known port" number of the server and a transient port number assigned to the client. In the exchange, the port numbers occur at two different positions, corresponding to *source port* and to *destination port*. Checking **either offset** with **AND** causes the analyzer to accept traffic passing between this pair of ports in either direction.

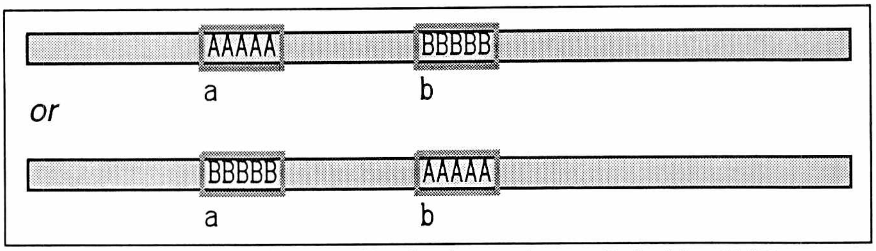

Figure 2–15. Effect of "either offset" when pairs are linked by AND.

In the unusual (but not impossible) case that you select OR as the relation and you also select **either offset**, any of the four frames shown in Figure 2–16 is a match. That filter would accept all traffic to or from port A and also all traffic to or from port B. That would include traffic between A and B, but it would also accept all their traffic with any other ports.

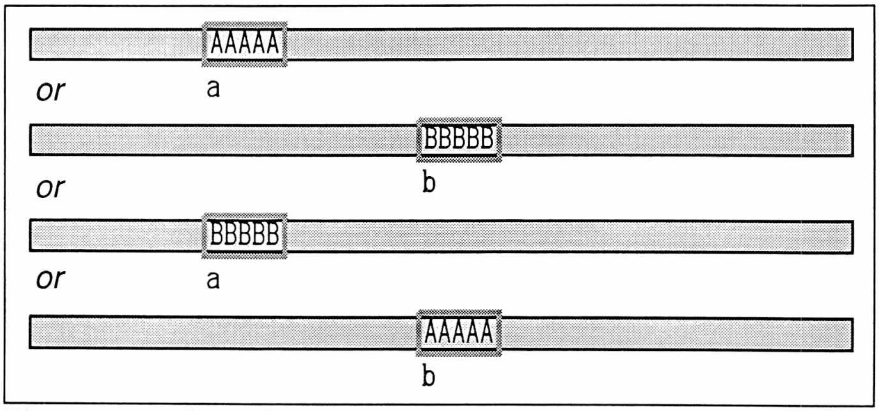

Figure 2–16. Effect of "Either offset" when pairs are linked by OR.

Effect of Match/Don't Match

For each of the four matches (called Match 1, Match 2, etc., or whatever other names you've given them), there is a radio control set to either **Match** or **Don't match**.

This setting reverses the evaluation of each of the patterns within a match. "Don't match" is evaluated *before* "either offset," and *before* the analyzer considers how AND or OR combine the pair to form a match. To see the effect of "Don't match," in Figure 2–13 through Figure 2–16, replace "AAAAA" by "Anything other than AAAAA" and replace "BBBBB" by "Anything other than BBBBB."

Examples

Suppose you describe a match as "pattern A at offset a." Then, setting "Don't match" means that (as far as this match is concerned) a frame should be accepted if, at offset a, it has something *other than* pattern A.

When a match contains a pair of patterns linked by AND (for example, "Pattern A at offset a" *and also* "Pattern B at offset b"), setting "Don't match" means that (as far as this match is concerned) a frame should be accepted if it contains something *other than* pattern A at offset a *and* something *other than* pattern B at offset b.

When you specify a pair of patterns linked by OR (for example, "Pattern A at offset a" *or* "Pattern B at offset b"), setting "Don't match" means that (as far as this match is concerned) a frame should be accepted if it contains something *other than* pattern A at offset a *or* something *other than* pattern B at offset b."

The Overall Outcome

Whether the frame is in fact accepted depends on the logic you've specified for combining the outcomes of the four matches (page 2–16).

Characters and Offset in an Individual Pattern

Each individual pattern is described by its position and the characters (or bits) it contains. A pattern may contain up to 32 characters.

Before you specify the pattern's content, first decide whether you will enter it as hexadecimal, as character or as binary. The default is hexadecimal. A set of radio controls at the bottom of the panel selects one of the three, as follows:

Hexadecimal Specify up to 16 bytes, each a pair of hex digits 00 through FF.

Character Specify up to 32 bytes, each an ASCII character enterable from the keyboard.

Binary Specify up to 4 bytes, as 32 binary bits each 0 or 1.

In hexadecimal or binary, an X means anything at that position. In ASCII character mode, Alt-x has the same effect; when you press Alt-x, the corresponding position is shaded, like this: ▨.

Data-Relative vs. Frame-Relative Offset

The position at which a set of characters is located is described by its *offset*. On token ring and certain other LANs, some frames contain a variable-length field called *source routing information.* The field, when present, comes after the DLC destination and source address, but before the regular data field. Thus, the position of data within a particular frame depends on whether that frame contains a source routing field. To allow for this uncertainty, you can state the offset in either of two ways:

Frame relative — Describe the location as the number of bytes from the start of the frame.

Data relative: — Describe the location as the number of bytes from the start of the frame's data segment; that is, from the start of the 802.2 frame data.

When a frame's source and destination are on the same network, it requires no forwarding, and the routing field is frequently (but not universally) omitted. When you are confident that all the frames that interest you are in the same format (all with a routing field or all without), it's safe to use a frame relative offset. Otherwise, you need a data relative offset. When you specify a data relative offset, the analyzer adjusts the offset to compensate for the length of each frame's routing field.

To enter a pattern for the Capture Filter (or for the Trigger or Display Filter)

1. Decide on the logic of your pattern. It can involve up to four matches, linked by logical operators AND or OR. (To review this part, see page 2–20 ff.)

   Each match can be a pattern-and-offset, or a pair of them. When the match contains a pair of patterns, the pair can be linked by AND or OR, with or without **either offset** (page 2–23).

   Each match can be a vote to include a frame when the data *match* or when the data *don't match* (page 2–24).

   If X stands for AND or OR (as appropriate), you can have

   (Match 1 X Match 2) X (Match 3 X Match 4)

2. In the main menu, highlight the option for which you are entering a pattern:
   - **Capture filter**, or
   - **Trigger**, or
   - **Display** and then **Filters**

   and from there highlight **Pattern match.**

3. Move the highlight to the match you will define (Match 1, 2, 3, or 4).

4. If you wish to name the match, press Enter. The analyzer opens a dialog box to receive your name for the match.

5. From the match name, move rightward to enter the match's specifications.

a. Indicate whether offsets for this pattern are:
   ▸ Frame relative
   Data relative (page 2–25)
   Highlight the desired entry and press Spacebar.

b. Indicate whether this match votes to include a frame on:
   ▸ Match
   Don't match
   Highlight the desired entry and press Spacebar.

c. If the match is for a pair of patterns, indicate by ✓ or X whether they should be accepted at either offset.
   Press Spacebar to toggle between ✓ and X.

d. Move the highlight down to the bottom of the panel to indicate the type of data for the pattern:
   ▸ Hexadecimal
   Character
   Binary
   Highlight the desired entry and press Spacebar.

6. Move the highlight to **Pattern=** and press Enter. The Sniffer analyzer opens a dialog box in which you can write the characters (Figure 2–17). Press Enter to record the pattern.

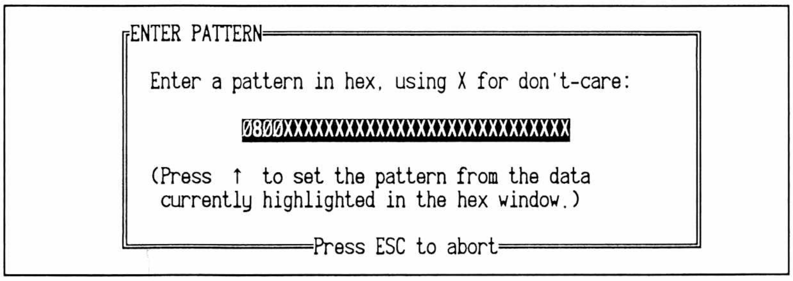

```
ENTER PATTERN

Enter a pattern in hex, using X for don't-care:

          Ø8ØØXXXXXXXXXXXXXXXXXXXXXXXXXXXX

(Press  ↑  to set the pattern from the data
 currently highlighted in the hex window.)

                 Press ESC to abort
```

Figure 2–17. Dialog box to enter a pattern.

7. Move the highlight to **offset=** and there press Enter. The Sniffer analyzer opens a box in which you can write the value of the offset. The offset is always in hexadecimal. Press Enter to record the offset.

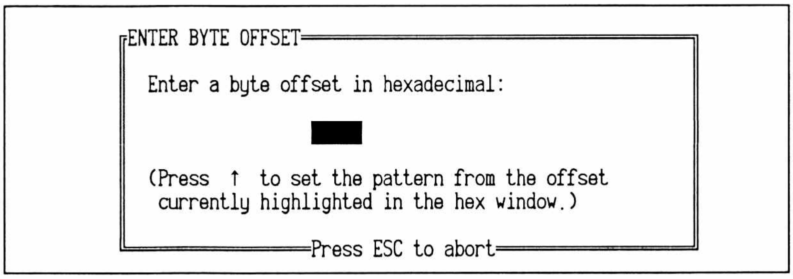

```
ENTER BYTE OFFSET

Enter a byte offset in hexadecimal:



(Press  ↑  to set the pattern from the offset
 currently highlighted in the hex window.)

                 Press ESC to abort
```

Figure 2–18. Dialog box to accept a pattern's offset.

8. Move the highlight leftward to the list of Matches. Select another match, and repeat Step 3. through Step 7. as necessary.

9. Move the highlight leftward to the list of Matches to set the relation between the separate matches.

   a. Set √ or X for each match to enable or disable it.

   b. Set the logical relation between Match 1 and Match 2, between Match 3 and Match 4, and between the two pairs of matches. For each:
      AND
      ➧ OR
      Highlight the desired entry and press Spacebar.

Cutting and Pasting a Pattern from the Display Hex Window

If you have already captured a frame that contains the pattern you want, you can have the analyzer copy its characters and its offset, so you don't have to type them. (This procedure makes use of the hex and detail views from the display menu, described in Chapter 5.)

To enter a pattern by copying and pasting

1. Set your display to include both the hex view and the detail view.

2. Find the frame that contains the pattern you want.

3. In the detail view, move the highlight to the field you want. This automatically moves the hex view's highlight to that field too.

4. While the field is highlighted, press F5 to return to the main menu.

5. Select **Capture filters**, then **Pattern match**, then to one of the four matches. Move rightward to specify its pattern.

6. If need be, move up to select **Frame relative** or **Data relative** offset.

7. Move to the line that contains **Pattern=** and press Enter. That brings you to the entry box shown earlier (Figure 2–17, page 2–27).

8. Press the Cursor Up key. That copies the characters that were highlighted in the hex view into the pattern entry dialog box. Press Enter to record the pattern.

9. Move to the line that contains **Offset=** and press Enter. That opens the other entry box. Press the Cursor Up key. That copies the pattern's offset into the offset dialog box. Press Enter to record the offset.

Example of Pattern Match in a Capture Filter

The following example is based on the use of Telnet over TCP/IP over Ethernet. However, the general strategy for setting a pattern match is not specific to the network or protocols of the example.

Suppose you experience a problem while using a terminal emulation package. While connected to a remote host, the network software confuses the actions of the Backspace and Delete keys, which (in this application) are supposed to do different things. Who is mixing them up? The emulator at the PC? The application at the host? The network software between them?

As a first step, you might examine what the emulator transmits to the host when you press Backspace and when you press Delete, as well as what the host echoes back to each. (Telnet often supports a terminal by transmitting one character at a time to the host. Usually, the host echoes each displayable character back to the terminal.)

To study what happens, you need a filter that accepts only those frames that include either Backspace or Delete embedded in a Telnet frame passing between the host and the terminal emulation program. You can't set a capture filter for the Telnet protocol explicitly (since you can't filter on high-level protocols during capture). But you can achieve the same result with pattern matching. Here's how you discover a pattern that identifies not just a Telnet frame, but one whose content involves the characters in question.

First set an address filter to select all frames sent from the terminal emulator. Then (from the terminal emulator) do something that involves Backspace and then something that involves Delete. Browsing through the frames thus captured, you can readily find Telnet frames addressed to the host. Examining them, you can see that each consists of:

- A DLC frame (with a DLC source and destination), and within that
- An IP frame (with an IP source and destination), and within that
- A TCP frame (with a TCP source and destination of "Telnet"), and within that
- A Telnet frame containing the record of a keystroke sent from the terminal emulator to the host.

To study how the program treats Backspace and Delete, you take advantage of the "cut and paste" facility to set a capture filter to match the following pattern. A frame is accepted when:

- It contains the IP protocol number for "TCP" (hex "06" at offset 17) and

- It contains the TCP code for Telnet data (indicated by a TCP source or destination port number of hexadecimal 17)

and

- Its IP source is the address of the PC running the terminal emulator, and its IP destination is the address of the host, or vice versa (by checking either offset)

and

- The Telnet data is either the code for delete (7F) or the code for backspace (08).

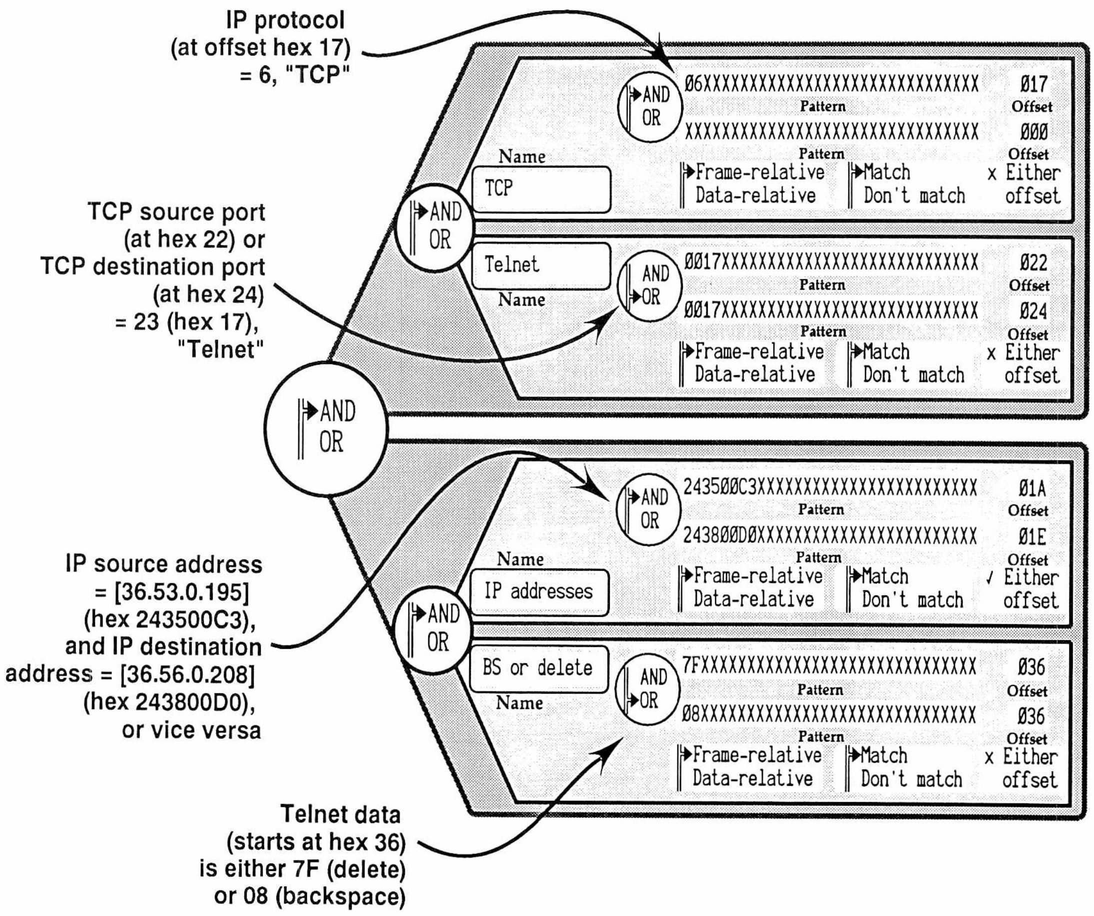

```
Name: TCP
Ø6XXXXXXXXXXXXXXXXXXXXXXXXXXXXXX   Ø17
Pattern                            Offset
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX   ØØØ
Pattern                            Offset
▶Frame-relative   ▶Match          x Either
 Data-relative     Don't match      offset

Name: Telnet
ØØ17XXXXXXXXXXXXXXXXXXXXXXXXXXXX   Ø22
Pattern                            Offset
ØØ17XXXXXXXXXXXXXXXXXXXXXXXXXXXX   Ø24
Pattern                            Offset
▶Frame-relative   ▶Match          x Either
 Data-relative     Don't match      offset

Name: IP addresses
2435ØØC3XXXXXXXXXXXXXXXXXXXXXXXX   Ø1A
Pattern                            Offset
2438ØØDØXXXXXXXXXXXXXXXXXXXXXXXX   Ø1E
Pattern                            Offset
▶Frame-relative   ▶Match          ✓ Either
 Data-relative     Don't match      offset

Name: BS or delete
7FXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX   Ø36
Pattern                            Offset
Ø8XXXXXXXXXXXXXXXXXXXXXXXXXXXXXX   Ø36
Pattern                            Offset
▶Frame-relative   ▶Match          x Either
 Data-relative     Don't match      offset
```

Figure 2–19. Example of a pattern match in the capture filter.

Figure 2–19 depicts the way pairs of individual patterns are combined to form matches, and the resulting matches are combined to form the filter. The diagram is shaded to emphasize the way the various components are grouped. The fields show data for the example just described. (Of course in a real situation the IP addresses would be different, but the values for "Telnet" and "TCP" are appropriate.)

Filters for Defective Frames

On a synchronous link or on Ethernet, the Sniffer analyzer can filter arriving frames for the presence or absence of certain defects. This filter is available only if the network interface card retains defective frames and passes them to the Sniffer server's CPU. The NIC for token ring simply discards defective frames. On token ring, these filters don't appear in your capture filters menu and you should skip this section of the manual.

The possible filters are indicated by check marks in Figure 2–20.

	WAN	Ethernet	
Good frames		✓	Frames having none of the defects listed below
Bad CRC	✓	✓	Frames having a defective CRC (cyclic redundancy check)
Fragments		✓	Frames having less than the minimum length. (These arise when a station detects that its transmission is colliding with the transmission of another station, and aborts its own transmission).
Bad alignment		✓	Frames whose length is not an integer multiple of 8 bits, and hence cannot be unambiguously resolved into bytes.

Figure 2–20. Filters for frame defects.

To set filters for frame defects

1. In the **Capture filters** menu, move the highlight downward to the last section.

   Result: You'll see the list of defect categories for the network the analyzer is monitoring. A ✓ mark indicates that frames in the category should be accepted, X indicates that they should be excluded.

2. Move the highlight to each category you wish to change. Press Spacebar to toggle between ✓ and X.

Tabulation of Defects when Categories Overlap

Defects are described by categories that overlap. A frame with bad alignment is probably a fragment and most likely has a bad CRC. To avoid counting the same defects several times, the Sniffer analyzer does not assign a frame to more than one category of defect. It checks for defects in a fixed sequence; when it finds a defect, it records that defect but does not count other defects in the same frame. The sequence is as follows:

1. Is the total length sufficient? If not, report that the frame is a fragment, and discontinue checking.
2. Is the alignment correct? If not, report that the frame has bad alignment and discontinue checking.
3. Is the CRC correct? If not, report that the frame has bad CRC.
4. If none of the above, report that the frame is a good frame.

Working Definition of "Fragment"

Ethernet requires that a frame contain at least 60 bytes. Any frame with fewer than 60 bytes is considered a fragment, and is so classified by the Sniffer analyzer.

Filters for Direction: From DTE or From DCE (WAN)

The synchronous link is bidirectional. However, you can elect to accept frames in one or the other direction or in both directions.

To filter frames by direction during capture from a WAN

1. From the main menu, move the highlight to **Capture filters,** and then to the panel to the right.
2. Move the highlight to the direction you wish to include or exclude:
   ✓ From DTE
   ✓ From DCE

   Press Spacebar to toggle between ✓ (include) and X (exclude).

Filters for Receiver Ready/Not Ready Frames (WAN)

In the process of setting up communication between the endpoints of a synchronous link, and to control the flow of frames once communication has been established, the devices (DTE and DCE) exchange frames designated RR ("receiver ready") and RNR. ("receiver not ready"). If you're investigating problems in handshaking or flow control these may be relevant. When your interest is primarily the higher level messages, the RR and RNR frames are usually irrelevant. The Sniffer analyzer provides filters that can include or exclude these frames. There are separate, independent filters for RR and RNR.

To filter RR or RNR frames during capture from a WAN

1. From the main menu, move the highlight to **Capture filters,** and then to the panel to the right.
2. Move the highlight to the type you wish to include or exclude:
   ✓ RR
   ✓ RNR

   Press Spacebar to toggle between ✓ (include) and X (exclude).

Trigger: Specifying When and How to Stop Capture

It's important to stop capture at the right moment. You want capture to stop while the frames that interest you are still in the capture buffer. When you've elected continuous capture, once the buffer fills, frames that arrived earlier are discarded to make room for new arrivals. So if you go on capturing too long, the frames you want may have been discarded.

The trigger mechanism gives you a way to stop capture.You can have capture halt immediately or somewhat later. Stopping immediately leaves you the trigger frame and frames that preceded it. Stopping later leaves you a buffer that also contains some frames that follow the trigger frame.

The Trigger Event

The analyzer stops capture when it detects a *trigger event.* In the distributed Sniffer system, the trigger event is the detection of a pattern in an accepted frame. [1]

When the trigger event occurs, the Sniffer analyzer posts the word TRIGGERED at the top left of the display screen. Then, or at a specified point thereafter, the server stops capturing frames and freezes the capture buffer. The capture buffer contains the *trigger frame*: that is, the frame that matched the trigger pattern.

At that point, the capture buffer also contains frames that preceded the trigger frame, or (optionally) frames that followed it. When you set up the trigger, you specify whether capture should cease immediately (so that the trigger frame is the last frame in the buffer),

[1] In the stand-alone Sniffer analyzer, the trigger event may also be a signal received at the analyzer's serial port.

or ceases when the trigger frame has progressed some percentage of the way through the capture buffer.

Setting the Trigger Pattern and its Position in the Buffer

There are two parts to setting the trigger:

- The pattern.
- The percentage of frames preceding the trigger event.

To set the trigger pattern

1. In the Sniffer analyzer's main menu, move the highlight to **Trigger**, and then to **Pattern**.
2. Set the pattern following the procedure already described for a pattern in the capture filter (see "Establishing a Complex Pattern by Combining Four Matches" on page 2–20 ff).

To set the position of the trigger frame in the capture buffer

1. In the Trigger menu, move the highlight up to the top item, **Stopping capture**.
2. Set the radio control to the appropriate percentage. To select one, highlight the desired option and press Spacebar. The choices are:

   Stop when full
   ▶ Continuous capture (the default)
   0% pretrigger
   25% pretrigger
   50% pretrigger
   75% pretrigger
   100% pretrigger

   (The value you pick determines the proportion of space in the capture buffer, not the number of frames.)

The effects of the various stopping options are summarized in Figure 2–21.

Option	Effect
Stop when full	Even if the trigger event has not occurred, capture stops when there is no more space in the capture buffer.
Continuous capture	The frames that arrived earlier are discarded to make room in the capture buffer for the newer arrivals. When the trigger event occurs, the analyzer (as usual) posts the word "Triggered" on the screen, but doesn't stop capturing. If nothing else happens to stop capture, sooner or later —depending on the size of the frames and the space in the buffer— arriving frames will displace those already captured and the trigger frame will be among those discarded.
0% pretrigger	Capture continues until the trigger frame is the oldest remaining in the capture buffer (frame 1), and all other frames follow it.
25% pretrigger	Capture continues until 25% of the space in the capture buffer is devoted to frames that arrived before the trigger frame.
50% pretrigger	As above, but 50% of the space in the capture buffer is devoted to frames that arrived before the trigger frame.
75% pretrigger	As above, but 75% of the space in the capture buffer is devoted to frames that arrived before the trigger frame.
100% pretrigger	Capture stops at once, so that the trigger frame is the last to arrive in the capture buffer and all other frames preceded it.

Figure 2–21. Effect of "stop capture" options in the trigger menu.

Examples of Trigger Patterns

Figure 2–22 shows two examples of the use of patterns to stop capture.

Token ring analyzer with IBM interpreter suite	Ethernet analyzer with XNS interpreter suite (monitoring a 3Com+ network)
To trigger on a frame reporting an SMB error, you might first set the capture filter to pass only NetBIOS frames. Then set the trigger to freeze the capture buffer when it finds that the primary SMB return code is not equal to 00 (zero means "normal"). In an IBM SMB frame, the primary return code is a single byte located at data-relative offset 27 (hex).	To trigger on a frame reporting an SMB error, you might first set the capture filter to pass only XNS frames. Then set the trigger to freeze the capture buffer when it finds that the primary SMB return code is not equal to 00 (zero means "normal"). In an XNS SMB frame, the primary return code is a single byte located at data relative offset 3D (hex).

Figure 2–22. Sample trigger pattern on token ring and Ethernet.

Marking the Trigger Frame

When the Sniffer analyzer matches the trigger pattern, it reports that fact by changing the label in the upper left corner of the screen from CAPTURING (Figure 2–31, page 2–45) to TRIGGERED (Figure 2–32). When you display the contents of the capture buffer, the trigger frame is marked with the letter T (see the section entitled Flags Option, page 5–20). During display, you can jump to the trigger frame, or if you wish, display time relative to the arrival of the trigger frame (see the section entitled Searching and Jumping, page 5--39).

The capture buffer never contains more than one frame marked T. Suppose you set **Stopping capture** so that capture continues after a trigger frame arrives. Suppose another matching frame arrives. What happens? Only the first is considered the trigger frame, and only the first is reported and marked.

The Capture Menu

After you've set up your capture filters and you've specified when capture should stop (in the trigger menu), there are still some options you may want to set in the capture menu before you press the button to start capture.

Buffer size

This item isn't really an option. It reports the number of kilobytes your machine has allocated to the capture buffer.

To display the size of the capture buffer

1. In the Capture menu, move the highlight upward to **Buffer=**.

   Result: The size of the buffer in kilobytes is visible as part of the item. If the buffer makes use of expanded memory, the number if kilobytes is followed by the letters EXP.

To display a summary of memory allocation

1. With **Buffer=** highlighted (as above), press Enter.

   Result: The analyzer displays a summary of memory utilization statistics, including DOS data space, expanded memory, and various components of the system heap (see Figure 2–23).

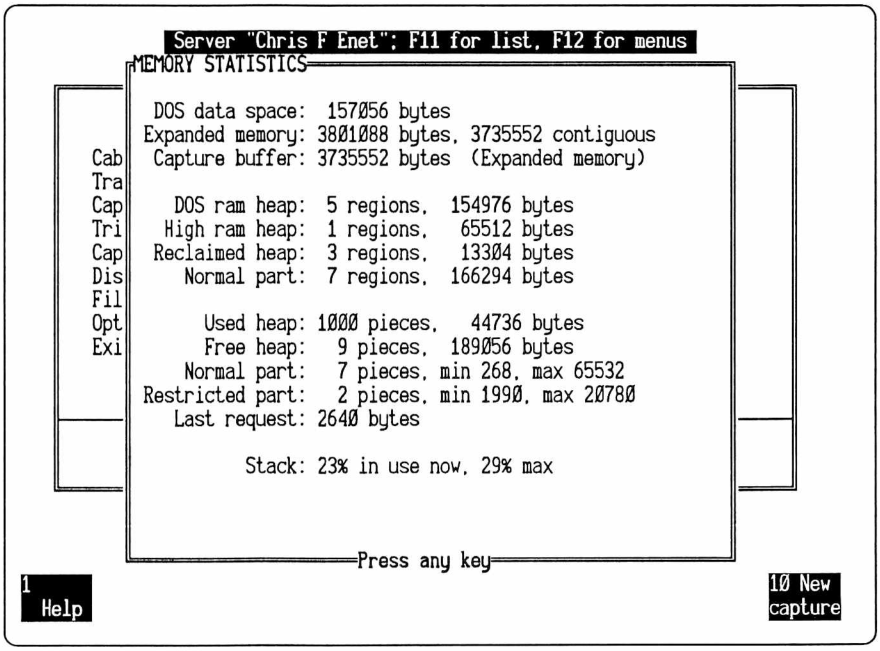

```
 Server "Chris F Enet": F11 for list, F12 for menus
MEMORY STATISTICS

     DOS data space:  157Ø56 bytes
    Expanded memory:  38Ø1Ø88 bytes, 3735552 contiguous
Cab  Capture buffer:  3735552 bytes  (Expanded memory)
Tra
Cap    DOS ram heap:  5 regions,  154976 bytes
Tri   High ram heap:  1 regions,   65512 bytes
Cap  Reclaimed heap:  3 regions,   133Ø4 bytes
Dis     Normal part:  7 regions,  166294 bytes
Fil
Opt       Used heap:  1ØØØ pieces,   44736 bytes
Exi       Free heap:    9 pieces,  189Ø56 bytes
        Normal part:    7 pieces, min 268, max 65532
    Restricted part:    2 pieces, min 199Ø, max 2Ø78Ø
       Last request:  264Ø bytes

              Stack:  23% in use now, 29% max

                     Press any key
1                                                    1Ø New
 Help                                                capture
```

Figure 2–23. Memory statistics display.

Frame size

When you're primarily interested in the frames' low-level headers, you can truncate frames that exceed a certain length and thereby fit more frames into the capture buffer. On a very busy network, truncation may also help avoid losing frames, since a longer frame takes slightly more time to store.

To limit capture to the first n bytes of a frame

1. In the **Capture** menu, move the highlight to **Frame size**, and then rightward to the panel of frame sizes.
2. Move the highlight up or down to the maximum length you want. Press Spacebar to move the radio control's selection arrow to the highlighted line. The choices are:

   32 bytes
   64 bytes
   128 bytes
   256 bytes
   512 bytes
   ▶ Whole frame (the default)

Effect of Truncating Frames

When each high-level frame is entirely contained within a lower-level frame, truncation leaves you the headers and discards part of the high-level data. Since the headers usually contain the information you need for analysis, little is lost by discarding the later parts of the frame.

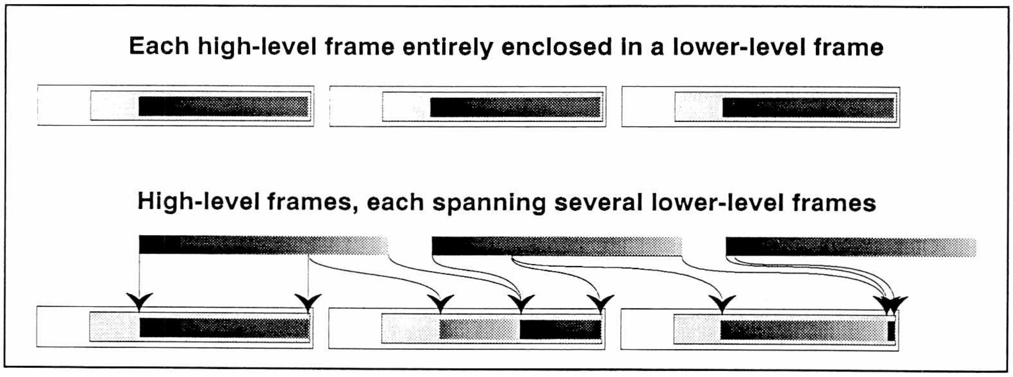
Each high-level frame entirely enclosed in a lower-level frame
High-level frames, each spanning several lower-level frames

Figure 2–24. Effect of high-level frames spanning multiple DLC frames.

However, some high-level protocols (for example, TCP) are byte-oriented rather than packet-oriented, and some protocols (for example, ISO or X Windows) permit very long messages. As a result a single ISO or X message may span several lower-level frames.

Figure 2–24 shows how a sequence of variable-length higher-level frames may be arbitrarily sliced and packed into frames of an intermediate byte-oriented protocol such as TCP. The start of a new spanned frame is not required to force a new lower-level frame. Thus, an X header (for example) may occur at any position in a TCP frame's data field. If your analysis requires keeping track of the headers of high-level spanned frames, it is essential to save whole frames, since the headers and boundaries of the highest levels may otherwise be lost.

Formatting the Screen Displayed During Capture

There are three displays that the analyzer generates as capture proceeds (and a fourth option to have no display). They're listed in Figure 2–25.

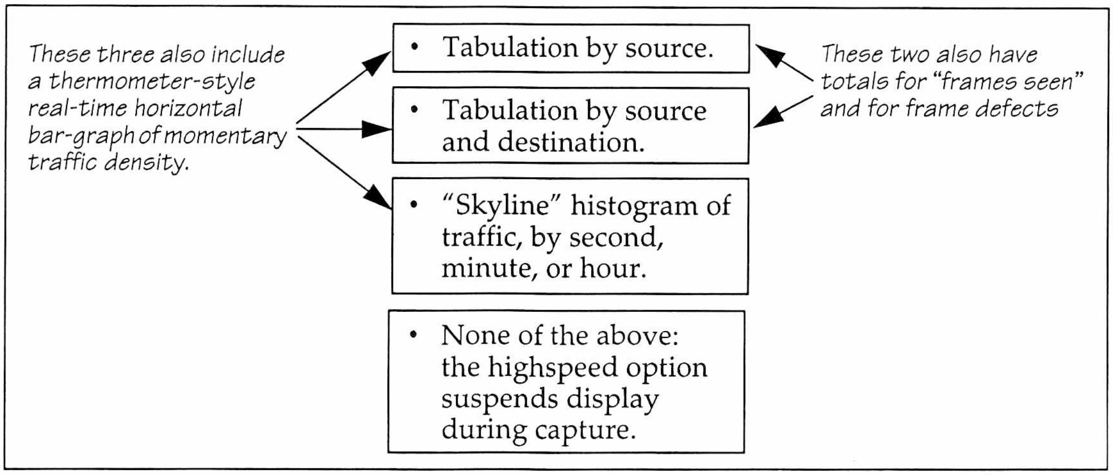

These three also include a thermometer-style real-time horizontal bar-graph of momentary traffic density.

- Tabulation by source.
- Tabulation by source and destination.
- "Skyline" histogram of traffic, by second, minute, or hour.
- None of the above: the highspeed option suspends display during capture.

These two also have totals for "frames seen" and for frame defects

Figure 2–25. Options for display during capture.

Options in the Capture menu select the form of display and the units. The options are:

Units for tables or skylines	▶Show frame counts Show Kbyte counts Show NW usage
Units for the "thermometer"	Linear bar scale ▶Log bar scale
Type of display	Individual counts ▶Pair counts Skylines

Examples of these displays are shown in detail later in this chapter; Figure 2–26 summarizes the places at which the choice of units and scale have their effect. The choices are described individually in the paragraphs that follow.

To select format of display during capture

1. In the **Capture** menu, move the highlight vertically to the desired options. At each radio control, press Spacebar to activate the one you want.

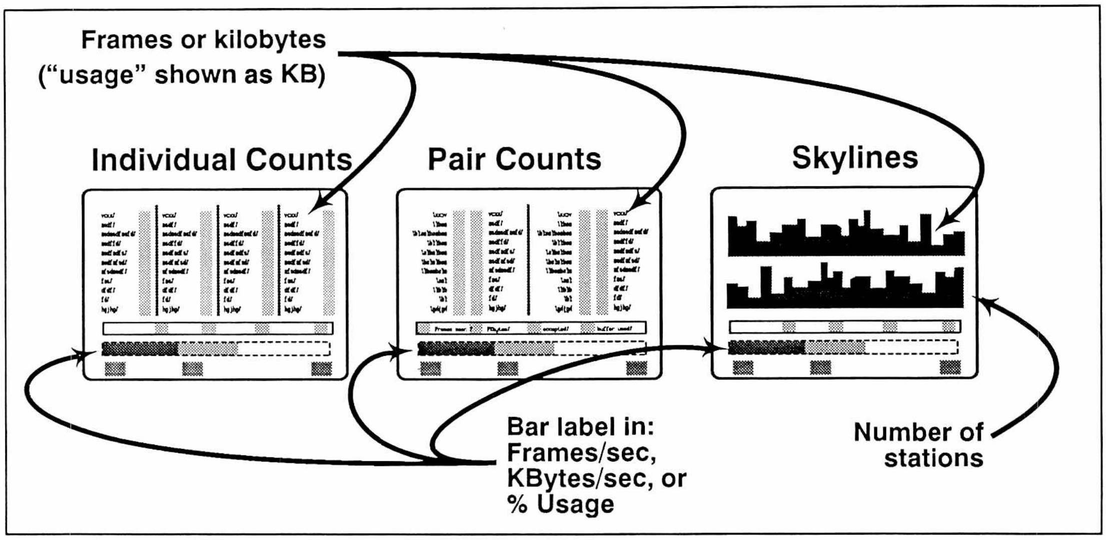

Figure 2–26. Effect of option on display during capture.

Units in Skylines, Tables, and the Traffic Density Bar Graph

You can select the units that the analyzer reports during capture. The choices and their effects are listed in Figure 2–27. The default is to count frames.

	Units in Counters (when used)	Units in Skylines (when used)	Units (on the traffic density bar graph)	Totals (above the traffic density bar graph)
Frame counts	Frames	Frames	Frames	unaffected (shows both frames and Kbytes)
Kilobytes	Kilobytes	Kilobytes	Kilobytes	
Network usage (LAN only)	Kilobytes	Kilobytes	Percentage of bandwidth	

Figure 2–27. Capture menu options for frames, kilobytes or usage.

Linear vs. Logarithmic Scale for the Traffic Density Bar Graph

You can select either a logarithmic or a linear horizontal scale for the traffic density bar graph. When the overall density is low, small variations are easier to see on a logarithmic scale (the default). These are your choices:

Linear	A fixed distance (say 1 cm.) corresponds to an absolute change in traffic density (say, 10,000 frames).
Logarithmic	A fixed distance (say 1 cm.) corresponds to a relative change (say, 10%),

Types of Display During Capture

This option selects the type of display during capture from among the following:

- Individual counts
- Pair counts
- Frame and call counts
- Skylines
- Highspeed (suppresses other displays).

The default is to display frame-type and call counts (on a synchronous link) or pair counts (on a LAN). Characteristics of the alternative displays are listed in Figure 2–28.

Display	Network	Description
Individual counts (source)	LAN only	Traffic is tabulated by DLC source address.
Pair counts (source and destination)	LAN only	Traffic is tabulated by DLC source and destination address
Frame-type and call count	WAN only	Traffic is tabulated by type, separately for DTE and DCE, and by logical call. The scrollable panel of logical calls may show all calls (the default) or only active calls.
Skylines	All	Histograms summarizing network activity at selectable intervals. The intervals are: every second (the default), every minute, or every hour.
	LAN	Two histograms: • Frames or kilobytes • Number of Stations
	WAN	Two histograms: • From DTE • From DCE
Highspeed	Ethernet only	All regular display is replaced by a single counter showing only the total of frames seen. Frames are stored in the interface card's temporary buffer, with about 400K capacity, instead of in normal capture buffer.

Figure 2–28. Characteristics of alternative displays during capture.

Items Always Displayed During Capture

Totals

The analyzer always reports the total number of frames it has seen (whether or not they passed its capture filters), and the total number of kilobytes they contained. On a synchronous link or token ring, these totals are reported directly as "frames seen" and "kilobytes seen." On other networks, there is no single total for "frames seen," but separate counts for various subtotals. For example, on Ethernet, there are totals for:

- Total number of good frames
- Total number of defective frames.
- Total number of lost frames.

Buffer utilization

As capture proceeds, the analyzer continuously updates a counter showing the percentage of the capture buffer that has been filled.

Momentary and "Highwater" Display of Traffic Density

As capture proceeds (whether with counters or "skylines") the Sniffer analyzer displays a thermometer-style horizontal bar to show real-time variations in traffic density (Figure 2–29). For LAN traffic, there's a single bar. On a synchronous link, there are two separate bars, one showing traffic from DTE and the other showing traffic from DCE. Each bar is updated several times a second. Each shows a moving average of the last half second's activity and the "high water"—that is, the maximum recorded during the current capture session.

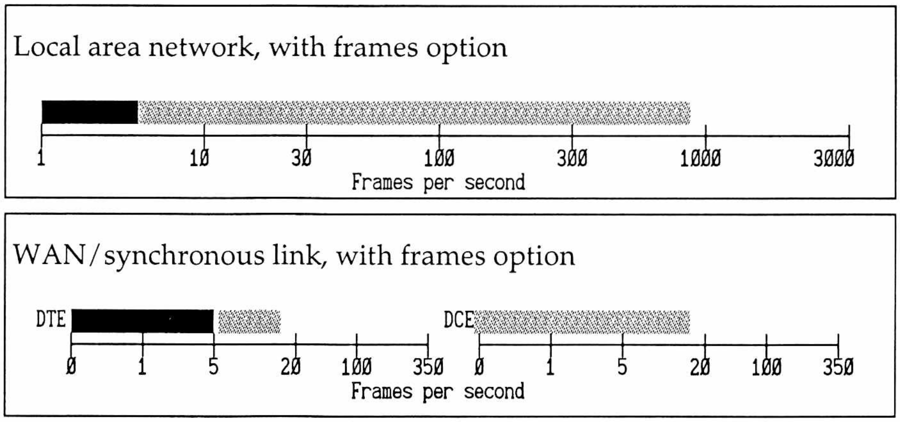
Local area network, with frames option

WAN/synchronous link, with frames option

Figure 2–29. Traffic density bar graph for LAN and for WAN.

Line Status on the Synchronous Link

Whether you choose counters or skylines, during capture the WAN/synchronous analyzer includes a one-line summary of the line's status (Figure 2–30). The display shows the RS232 indicators RxC, TxC, RxD, TxD, CTS, RTS, DSR and DTR. The condition of each is indicated by the symbol ↑ ↓ or ↕.

Interpreting the Arrows

The symbol ↑ indicates a logical 1, ↓ indicates a logical 0, and ↕ indicates that the indicator's status has changed in the last second. But what does a logical 1 mean? That depends on the way the server is connected to the synchronous line. When the connection uses the DB-25 (RS232) connector supplied with the server, 1 (the up arrow) indicates active, enabled, or OK. But when the connection uses the optional V.35 interface pod, the significance of the signals for CTS, RTS, DSR and DTR is reversed: for those indicators, 0 (the down arrow) indicates *active, enabled* or *OK*. See the section entitled "Connecting a WAN Server" in Chapter 3 of *Distributed Sniffer System: Installation and Operations Manual.*

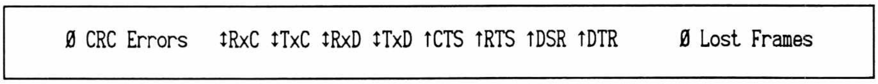

```
Ø CRC Errors   ↕RxC ↕TxC ↕RxD ↕TxD ↑CTS ↑RTS ↑DSR ↑DTR      Ø Lost Frames
```

Figure 2–30. Line status indicators on a synchronous link.

What the Capture Screens Look Like

The examples that follow illustrate the various types of display during capture —counters and histograms— on both local and wide area networks.

Capture with Pair Counts on a LAN

The analyzer allocates a slot for each pair of communicating stations, and updates them as frames are accepted. The counters show either frames or kilobytes, as you elected in the **Capture** menu.

The table is initially blank. As the server notices traffic between a station pair that it hasn't seen before, it adds an entry. For each pair, there's a counter for traffic in the direction first observed, and a second for traffic in the reverse direction. The number closest to the station's name describes transmissions *from* that station *to* the other station.

The analyzer adds pairs as it encounters them until it has used all available positions on the screen. After that, it continues to update the grand totals and the counts for the entries on the screen, but does not record details of pairs that aren't on the screen. (To clear the screen and start the tabulation anew, press F4. This doesn't affect the frames in the capture buffer.)

Figure 2–31 shows the screen following a lengthy capture with pair counts. Although it was recorded from an Ethernet LAN, the screen would be quite similar for token ring. Figure 2–33 (page 2–47) shows a screen reporting capture on a synchronous link.

In Figure 2–31, most stations are identified by names because a name table containing most of the DLC addresses had been created before capture started. The table had no name for the station shown as Exceln923931. The address that appears as AAAAAAAAAAAA (in both source and destination!) is spurious; in this particular case, it appeared because a station with a defective card transmitted several fragments when first turned on. Each fragment contained nothing but the character hex A in every field. Setting the capture filter to exclude defective frames would eliminate this spurious entry.

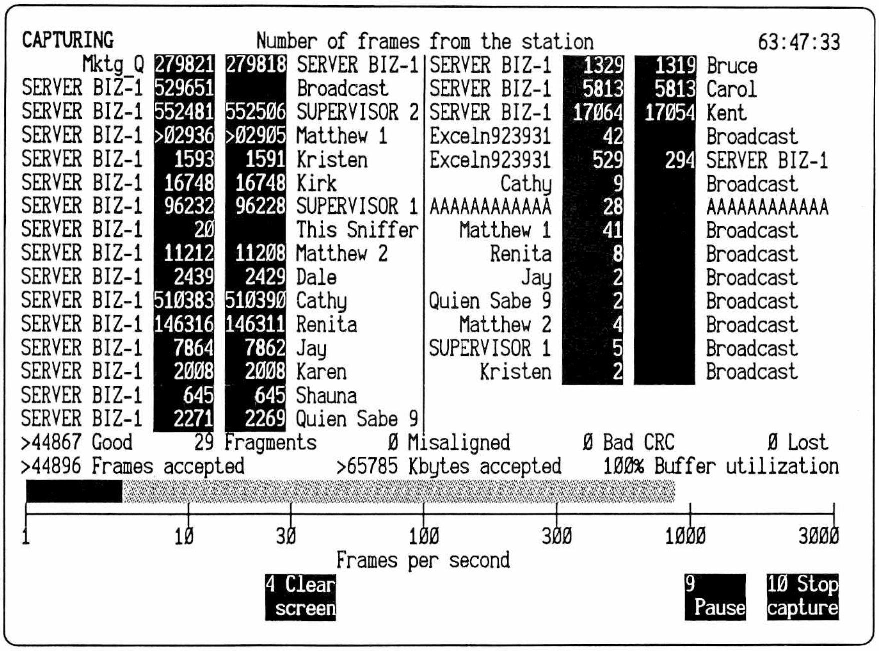

```
CAPTURING             Number of frames from the station              63:47:33
      Mktg_Q 279821 279818 SERVER BIZ-1|SERVER BIZ-1  1329  1319 Bruce
SERVER BIZ-1 529651        Broadcast   |SERVER BIZ-1  5813  5813 Carol
SERVER BIZ-1 552481 552506 SUPERVISOR 2|SERVER BIZ-1 17064 17054 Kent
SERVER BIZ-1 >02936 >02905 Matthew 1   |Exceln923931    42       Broadcast
SERVER BIZ-1   1593   1591 Kristen     |Exceln923931   529   294 SERVER BIZ-1
SERVER BIZ-1  16748  16748 Kirk        |       Cathy     9       Broadcast
SERVER BIZ-1  96232  96228 SUPERVISOR 1|AAAAAAAAAAAA    28       AAAAAAAAAAAA
SERVER BIZ-1     20        This Sniffer|   Matthew 1    41       Broadcast
SERVER BIZ-1  11212  11208 Matthew 2   |      Renita     8       Broadcast
SERVER BIZ-1   2439   2429 Dale        |         Jay     2       Broadcast
SERVER BIZ-1 510383 510390 Cathy       |Quien Sabe 9     2       Broadcast
SERVER BIZ-1 146316 146311 Renita      |   Matthew 2     4       Broadcast
SERVER BIZ-1   7864   7862 Jay         |SUPERVISOR 1     5       Broadcast
SERVER BIZ-1   2008   2008 Karen       |     Kristen     2       Broadcast
SERVER BIZ-1    645    645 Shauna      |
SERVER BIZ-1   2271   2269 Quien Sabe 9|
>44867 Good        29 Fragments       0 Misaligned       0 Bad CRC        0 Lost
>44896 Frames accepted       >65785 Kbytes accepted     100% Buffer utilization

1           10       30          100         300         1000        3000
                          Frames per second
                  4 Clear                                         9      10 Stop
                   screen                                   Pause  capture
```

Figure 2–31. Capture screen showing pair counts.

Capture With Individual Counts on a LAN

Individual counts (Figure 2–32) work in much the same way as pair counts. The analyzer adds an entry to the screen for each station that transmits, in the order encountered. Because these entries record the sender but not the destination, fewer total entries are required, and each takes less space. So the screen has room for a greater number of entries. As with pair counts, the server adds entries until all positions on the screen are filled. Then it continues to update the grand totals and the counts for the entries on screen, but does not record details of stations that aren't on screen. (To clear the screen and start tabulation anew, press F4; that doesn't affect frames in the capture buffer.)

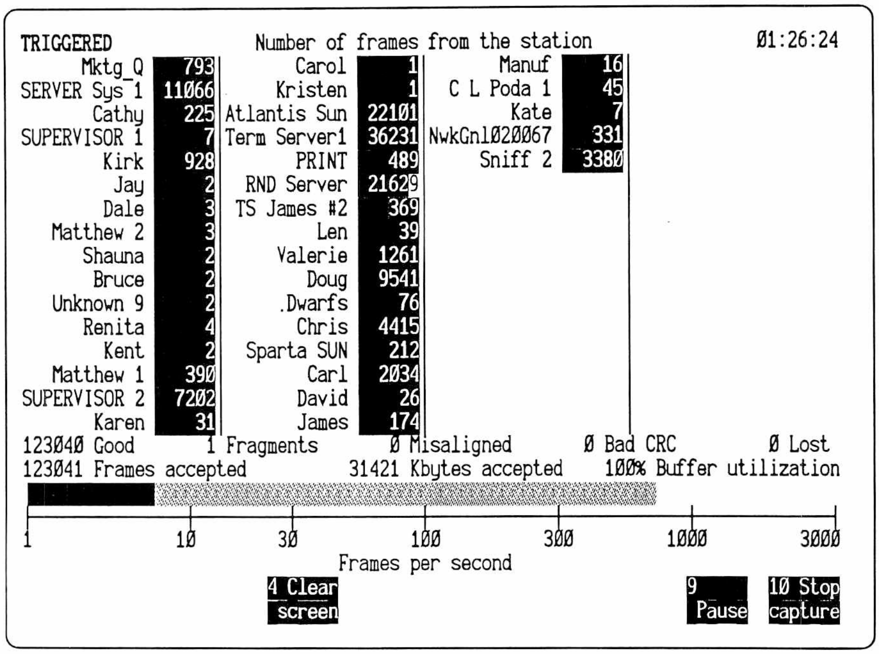

```
TRIGGERED                    Number of frames from the station                 Ø1:26:24
      Mktg_Q    793        Carol      1       Manuf    16
SERVER Sys 1  11Ø66      Kristen      1   C L Poda 1   45
       Cathy    225 Atlantis Sun  221Ø1         Kate    7
SUPERVISOR 1      7 Term Server1  36231 NwkGn1Ø2ØØ67  331
        Kirk    928        PRINT    489      Sniff 2 338Ø
         Jay      2   RND Server  21629
        Dale      3  TS James #2    369
   Matthew 2      3          Len     39
      Shauna      2      Valerie   1261
       Bruce      2         Doug   9541
   Unknown 9      2       .Dwarfs    76
      Renita      4        Chris   4415
        Kent      2   Sparta SUN    212
   Matthew 1    39Ø         Carl   2Ø34
SUPERVISOR 2   72Ø2        David     26
       Karen     31        James    174
12Ø4Ø Good          1 Fragments        Ø Misaligned        Ø Bad CRC          Ø Lost
123Ø41 Frames accepted          31421 Kbytes accepted     1ØØ% Buffer utilization
```

Figure 2–32. LAN capture screen showing individual counts.

Capture with Frame Counts on a Synchronous WAN

Counters

On a WAN/synchronous link, the metering screen during capture is divided into three zones (see Figure 2–33). On the left there are counters showing either frames or kilobytes for each of twelve HDLC types, totaled separately by direction (from DTE and from DCE).

X.25 or SNA

The center zone has counters showing either frames or kilobytes at the next protocol level, either X.25 or SNA (according to your choice in the frame type item in the options menu). Since only information frames have SNA or X.25 content, the total number of frames in the center panel may be lower than the total in the left panel.

Logical calls

In the third zone, the analyzer builds a table of logical calls, in the order encountered in the traffic. Each call is identified by its LCN and by its call address (if known). The LCN is composed of a logical channel group number and logical channel number.

A column at the right shows whether each call originated from DTE or from DCE. For each call, the analyzer tabulates traffic on that connection in each direction.

The LCN table is open ended, depending on the number of calls identified while capture continues. When the number of calls becomes larger than the number of lines allocated in the viewing panel, you can scroll the list by pressing the Cursor Up or Cursor Down keys, or F7 or F8.

Active vs. completed calls

The logical call table includes both calls that are still active and calls that have been completed. The counts for calls that are still active are highlighted. (In Figure 2–33, the counts 79 and 60 for LCN 014 are highlighted, indicating that the call is active, whereas those on the preceding and following rows are no longer active and hence are not highlighted.)

You may restrict the display to completed calls by pressing F2. Pressing F2 again restores the display to include all calls. Once a call has been completed, its logical call number may be reused, so the inactive calls may include multiple instances of the same LCN.

Unknown LCN Address

A logical call's address is contained in its first frame. If a call began before the Sniffer analyzer started to capture, its LCN is known but not its address. The address for such a call is shown as a row of dashes.

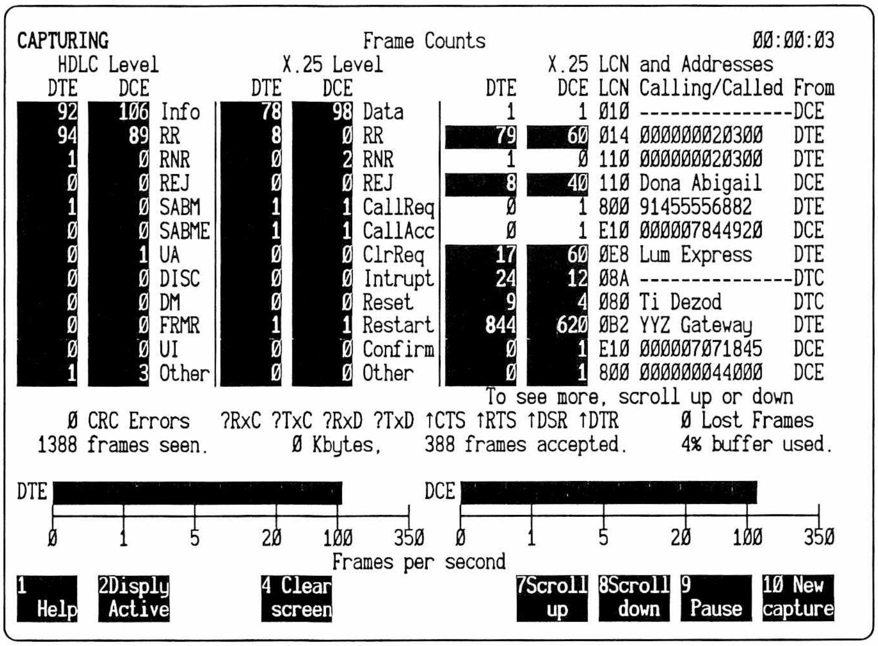

```
CAPTURING                        Frame Counts                           ØØ:ØØ:Ø3
  HDLC Level            X.25 Level                   X.25 LCN and Addresses
 DTE   DCE            DTE   DCE                DTE    DCE LCN Calling/Called From
  92   1Ø6 Info    |   78    98 Data     |       1      1 Ø1Ø ----------------DCE
  94    89 RR      |    8     Ø RR       |      79     6Ø Ø14 ØØØØØØØ2Ø3ØØ   DTE
   1     Ø RNR     |    Ø     2 RNR      |       1      Ø 11Ø ØØØØØØØ2Ø3ØØ   DTE
   Ø     Ø REJ     |    Ø     Ø REJ      |       8     4Ø 11Ø Dona Abigail   DCE
   1     Ø SABM    |    1     1 CallReq  |       Ø      1 8ØØ 91455556882    DTE
   Ø     Ø SABME   |    1     1 CallAcc  |       Ø      1 E1Ø ØØØØØ784492Ø   DCE
   Ø     1 UA      |    Ø     Ø ClrReq   |      17     6Ø ØE8 Lum Express    DTE
   Ø     Ø DISC    |    Ø     Ø Intrupt  |      24     12 Ø8A ----------------DTC
   Ø     Ø DM      |    Ø     Ø Reset    |       9      4 Ø8Ø Ti Dezod       DTC
   Ø     Ø FRMR    |    1     1 Restart  |     844    62Ø ØB2 YYZ Gateway    DTE
   Ø     Ø UI      |    Ø     Ø Confirm  |       Ø      1 E1Ø ØØØØØ7Ø71845   DCE
   1     3 Other   |    Ø     Ø Other    |       Ø      1 8ØØ ØØØØØØØ44ØØØ   DCE
                                            To see more, scroll up or down
    Ø CRC Errors   ?RxC ?TxC ?RxD ?TxD ↑CTS ↑RTS ↑DSR ↑DTR       Ø Lost Frames
 1388 frames seen.          Ø Kbytes,    388 frames accepted.     4% buffer used.

DTE                                         DCE

 Ø     1     5     2Ø    1ØØ   35Ø          Ø     1     5     2Ø    1ØØ   35Ø
                             Frames per second
1        2Disply      4 Clear                   7Scroll 8Scroll 9       1Ø New
  Help    Active       screen                      up     down   Pause  capture
```

Figure 2–33. WAN capture, with counts for HDLC, X.25 and logical call.

Names for Addresses

You can supply names for the source or destination of a logical call. The mechanism is the same as names for station addresses on a LAN (see page 5–12). The analyzer substitutes the name for the remote address (the source on a call from DCE, the destination on a call from DTE). Names for addresses are visible in the right panel of Figure 2–33.

Capture with Skylines

A skyline is a real-time graph of traffic density. It shows a histogram that is updated by adding a new column at the right every second, minute, or hour, depending on the time scale you selected before you started capture. With the skyline (as with the table of counters), you also get a count of total frames, but no other itemization.

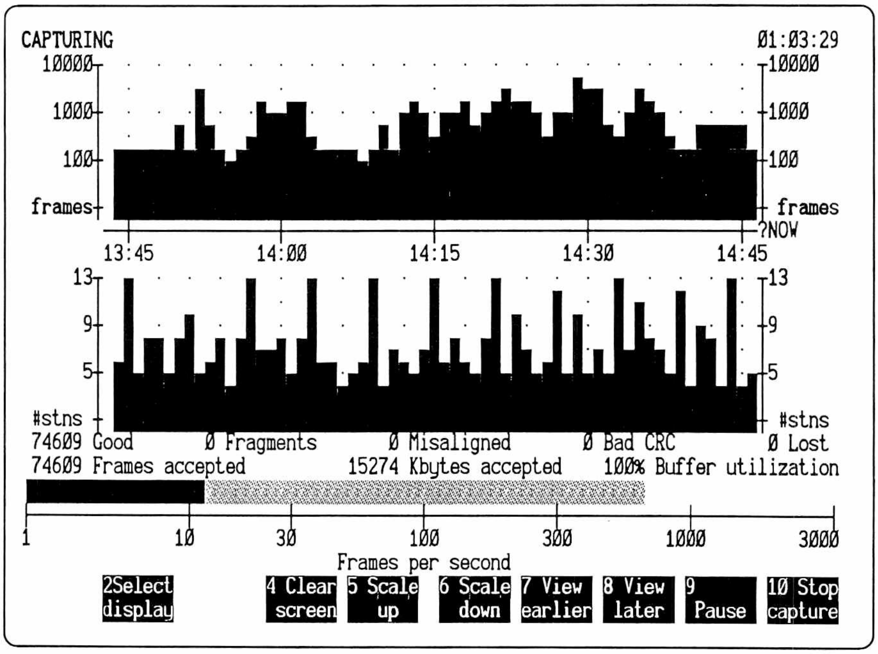

```
CAPTURING                                                        Ø1:Ø3:29
74609 Good        Ø Fragments        Ø Misaligned       Ø Bad CRC      Ø Lost
74609 Frames accepted          15274 Kbytes accepted    1ØØ% Buffer utilization
```

Figure 2–34. Skyline display on a local area network

A skyline always has two graphs, one above the other, with a common horizontal time scale. On a LAN (Figure 2–34), the upper histogram shows either the number of frames or the number of kilobytes transmitted during the interval, according to your choice of units in the capture menu. The lower histogram always shows the number of stations active during the interval.

On a WAN/synchronous link (Figure 2–35), the upper histogram shows traffic from DTE, while the lower histogram shows traffic from DCE.

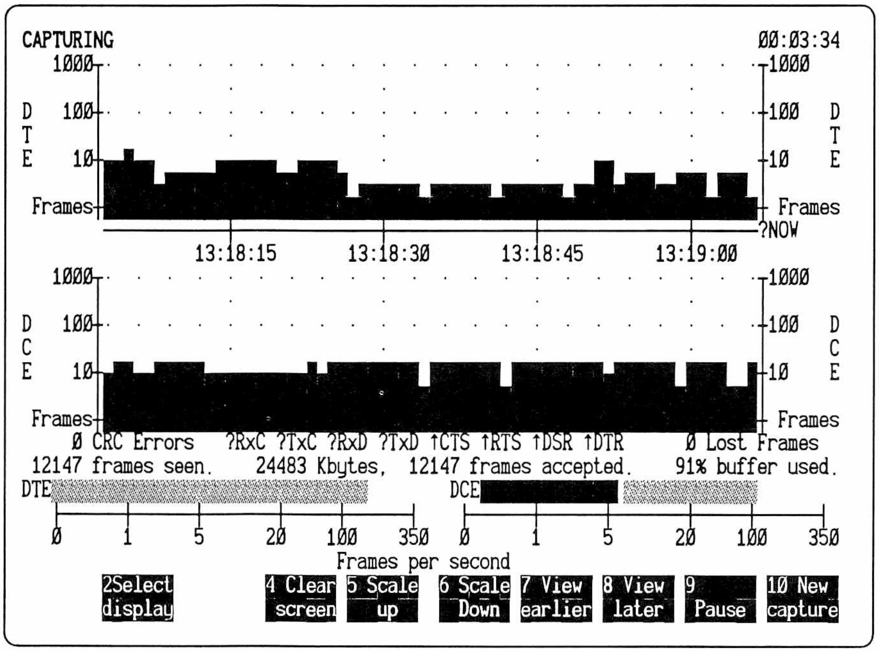

```
CAPTURING                                                    ØØ:Ø3:34
Ø CRC Errors   ?RxC ?TxC ?RxD ?TxD ↑CTS ↑RTS ↑DSR ↑DTR    Ø Lost Frames
12147 frames seen.   24483 Kbytes,  12147 frames accepted.  91% buffer used.
```

Figure 2–35. Skyline display on a WAN synchronous link.

Vertical Scale and Horizontal Time Span of the Skylines

You can have the vertical axes calibrated on either a linear or a logarithmic scale, as you request in the capture menu. Whichever you choose applies to all the histograms.

You can also adjust the scale of each vertical axis independently. Pressing F5 (labeled scale up) gives you taller bars (on a smaller range), while pressing F6 does the reverse. You can press either of these while collection proceeds.

The scale up or scale down keys affect only one of the skylines at a time. The one affected is the one for which the label on the vertical axis is highlighted. In Figure 2–34 or Figure 2–35, the word "frames" that labels the upper skyline appears in bold (or in a contrasting color on an external color monitor) to indicate that the upper skyline is active and will respond to scale up or scale down, whereas on the lower skyline the label #stns is not.

To move the highlight (and the effect of F5 and F6) from one skyline to the other, press F2 or the Tab key.

Capture Source: Live or Playback

The source from which frames are captured is controlled by the capture option labeled **From <xxxxx>**. At the bottom of the capture menu, you'll see **From** followed by the name of the network, for example **From Ethernet** or **From token ring**, etc., as appropriate.

When you capture from a file, this part of the menu shows **From** followed by the name of the file.

To capture from a file

1. From the analyzer's main menu, move the highlight to **Capture**.
2. Move the highlight to **From <xxx>** and press Enter. (**From** is at the bottom of the menu; initially, it may be below the lower edge of the panel. Scroll down until you come to it.)

   Result: The Sniffer analyzer opens a dialog box from which you can select the source you want. Initially, the dialog box shows the names of trace files (that is, files of saved frames) in the CAPTURE directory of the Sniffer server's hard drive.
3. Within the dialog box, move the highlight up or down to select the file you want, and press Enter.

   Result: The analyzer closes the dialog box, and inserts the name of the selected file in the capture source.
4. To move to a subdirectory within the current directory: highlight its name and press Enter.
5. To move to a directory closer to the root: move the highlight upward to the first entry (labeled " . .") and press Enter.

To capture from the network

... when the current capture source is a file

1. Proceed as described in "To capture from a file."
2. Within the dialog box, move the highlight to the first line. It contains the name of the network (for example, <Ethernet>). Press Enter.

   Result: The analyzer closes the dialog box, and inserts the name of the selected network in the capture source.

Highspeed Capture

Under certain conditions, sustained highspeed traffic might outrun the analyzer's ability to capture every frame. On Ethernet, to permit the server to give maximum time to the process of capture, the **highspeed** option stores frames in the interface card's temporary buffers instead of in the area of main storage dedicated to the capture buffer. This speeds up processing considerably. However, it also:

- Limits storage to the approximately 400KB of the NIC's buffer,
- Bypasses filtering, and

- Prevents the analyzer from generating its usual display of counters or skylines.

You still get a count of total frames, displayed in a small panel superimposed on the usual display, which is otherwise frozen.

The Capture menu includes the **highspeed** option only on Ethernet (the only network for which the interface card provides programmable access to the local buffer).

To start capture in highspeed mode on Ethernet

1. In the main menu, highlight **Capture** and then move rightward to its submenu. Move the highlight down to **Highspeed** (it's at the bottom, initially out of sight below the lower edge). Press Enter to toggle X (inactive) to √ (active).

2. In the **Capture** menu, select **Individual counts** or **Pair counts**, as you prefer. (You can select **Skylines**, but if you do, during highspeed capture the effect is the same as selecting **Pair counts**.)

3. Press F10 to start capture.

   Result: The analyzer puts up the framework for the counts you requested but does not fill in any data. A rectangle in the center is superimposed for the highspeed counts. As capture proceeds, only the central rectangle is updated.

   When you press F10 (Stop capture) or F9 (Pause), the analyzer transfers the frames that have accumulated in the NIC's buffer to the analyzer's capture buffer. Then it displays the accumulated statistics in the format you requested (pair counts or individual counts).

Figure 2–36 illustrates a screen visible during highspeed capture.

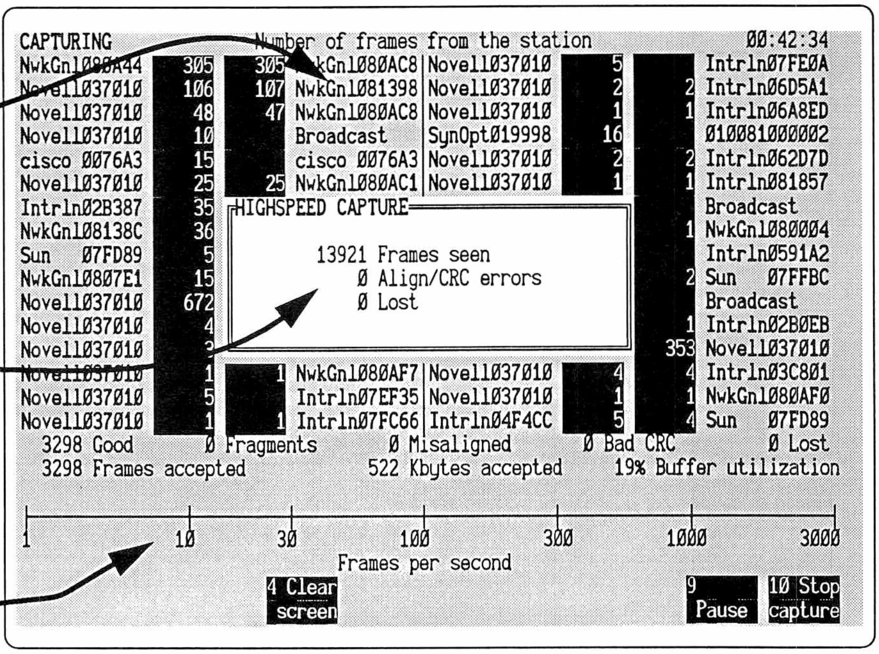

```
CAPTURING              Number of frames from the station                00:42:34
NwkGn1080A44  305   305 NwkGn1080AC8|Novell037010     5                Intrln07FE0A
Novell037010  106   107 NwkGn1081398|Novell037010     2              2 Intrln06D5A1
Novell037010   48    47 NwkGn1080AC8|Novell037010     1              1 Intrln06A8ED
Novell037010   10       Broadcast   |SynOpt019998    16                010081000002
cisco 0076A3   15       cisco 0076A3|Novell037010     2              2 Intrln062D7D
Novell037010   25    25 NwkGn1080AC1|Novell037010     1              1 Intrln081857
Intrln02B387   35  ┌HIGHSPEED CAPTURE══════════════════════┐           Broadcast
NwkGn108138C   36  │                                       │         1 NwkGn1080004
Sun   07FD89    5  │        13921 Frames seen              │           Intrln0591A2
NwkGn10807E1   15  │            0 Align/CRC errors         │         2 Sun   07FFBC
Novell037010  672  │            0 Lost                     │           Broadcast
Novell037010    4  │                                       │         1 Intrln02B0EB
Novell037010    3  └═══════════════════════════════════════┘       353 Novell037010
Novell037010    1     1 NwkGn1080AF7|Novell037010     4              4 Intrln03C801
Novell037010    5       Intrln07EF35|Novell037010     1              1 NwkGn1080AF0
Novell037010    1     1 Intrln07FC66|Intrln04F4CC     5              4 Sun   07FD89
  3298 Good          0 Fragments        0 Misaligned       0 Bad CRC          0 Lost
  3298 Frames accepted               522 Kbytes accepted     19% Buffer utilization

 |           |         |            |             |             |            |
 1           10        30           100           300           1000         3000
                               Frames per second
                4 Clear                                         9       10 Stop
                 screen                                          Pause  capture
```

Individual counts or pair counts are frozen until you pause or stop capture

During highspeed capture, only the counters in the center overlay are updated!

Frame counts and bar graph are also frozen until you pause or stop capture

Figure 2–36. Screen during highspeed capture on Ethernet.

Starting Capture

When you have set up the capture filters and triggers and have specified the capture options you want, you're ready to start capture.

To start capture

1. Check that you've set the capture filters and options as you want them.

2. Press F10 (New Capture).

   Or, in the main menu, move the highlight to **Capture** and press Enter.

   Result: The Sniffer analyzer starts capturing frames. The screen changes to the type of display you've chosen (frame counts, pair counts, individual counts, or skylines) with the scales and measures you indicated.

Ethernet Cable Test

On Ethernet only, if this is the *first* capture since the analyzer program started, the analyzer runs a *cable test* before it starts capturing. In the event it detects a cable fault, it displays the message "`There appears to be a cable or transceiver fault`," followed by some diagnostic information. You can then stop to investigate or proceed with capture. The cable tester is described in Chapter 3.

Stopping Capture

Once started, capture continues until one of the following happens:

- You press F10 to stop capture;
- The trigger event occurs (if you specified that capture should stop then);
- The buffer fills (if you specified that capture should stop then);
- You press F9 to pause capture. That permits you to adjust the display or capture filters and then resume capturing frames.

What You Can Do While Capture is Paused

If you press F9, the analyzer pauses its capturing. Frames are no longer being captured, but the frames already captured remain in the capture buffer. At that point, you can press function keys as follows:

F1 **Help** This is the same help facility available throughout the analyzer (but not accessible while actively capturing).

F9 **Resume** When you press **Resume**, capture continues. The next frame to be captured is appended to the capture buffer immediately following the last frame captured before you press **Pause**. There is no special mark to show that a pause occured. Frames that arrived during the pause are simply missing.

F4 **Clear screen** This clears the current graph or tabulation on the screen, but doesn't affect frames in the capture buffer.

F6 **Capture options** This permits you to change the capture options, but doesn't affect frames in the capture buffer.

If you make any changes to the options, filters, etc., when you press **Return** (F6) or **Resume** (F9), the server will clear the screen before resuming.

F3 **Data display** This stops capture and takes you to display; after this, you can't resume your earlier capture (but you can start a new capture, which will first clear the capture buffer).

What You Can Do When Capture Has Stopped

Once capture has stopped, you can do any of the following with the frames in the capture buffer:

Display them

Chapter 5 tells you how to display frames once they're in the capture buffer.

Save them

The saved frames go to a file on the server's hard disk. (You may subsequently move that file to the console, using the file transfer utility).

A file of captured frames can be used as:

- A source from which you "capture" frames (see page 2–49).
- A source from which you load the capture buffer (see page 5–68).

Discard them

When you start a new capture or exit from the Sniffer analyzer program without having saved the contents of the capture buffer, the analyzer warns you that if you proceed it will discard the frames now in the buffer.

CHAPTER THREE: NETWORK–SPECIFIC TESTING 3

Chapter 3. Network-Specific Testing

Chapter Overview

Chapter 3 describes analyzer services that are specific to a particular network. Such a service shows up in the analyzer's main menu only when you're using a server for that network. Since most of the analyzer's functions are available across the board to any of the supported networks, this is a short chapter. In the current version of the distributed Sniffer system, there is only one network-specific service: the Ethernet cable tester.

This chapter describes:

- What the cable tester does
- How you invoke it
- How you interpret its reports of a short or an open circuit.

The chapter concludes with some notes on limitations of the cable tester.

Ethernet Cable Tester

On Ethernet, the Sniffer analyzer provides a cable tester: a means to check and report cable faults. The analyzer's main menu includes an option labeled **Cable tester** (Figure 3–1). The cable test uses the server's "monitor" interface card as a time-domain reflectometer. During the test, the analyzer repeatedly emits a pulse and listens for the echo characteristic of certain types of faults.

To test the Ethernet segment for cable faults

1. In the analyzer's main menu, move the highlight to **Cable Tester** and press Enter (Figure 3–1).

   Result: The analyzer repeatedly emits a test signal on the Ethernet segment to which the server's "monitor" card is attached. The analyzer overlays a display reporting what faults —if any— have been detected, and updates it as the test is repeated.

2. To terminate the test, press Esc.

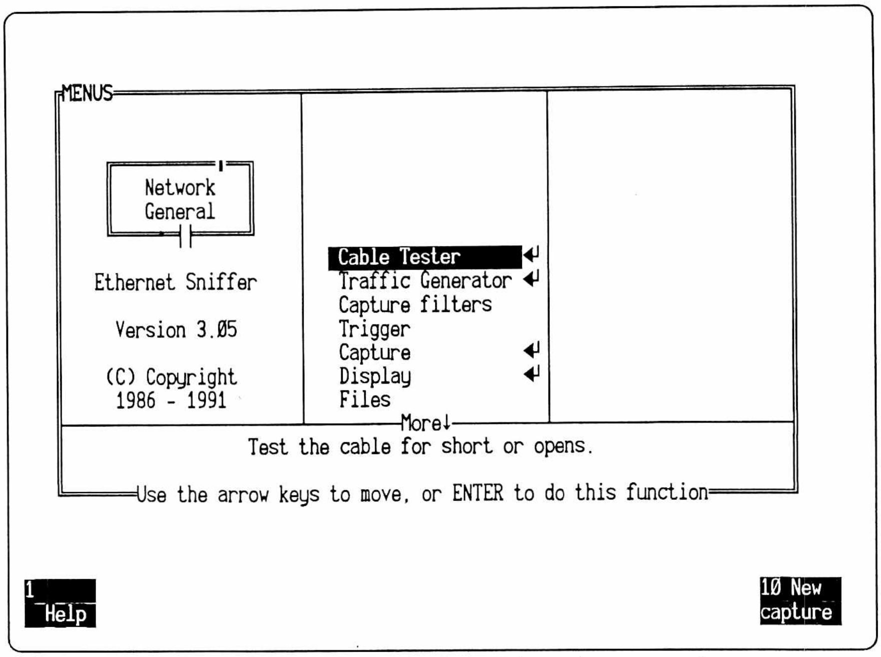

```
MENUS
   Network
   General

                        Cable Tester       ↵
 Ethernet Sniffer       Traffic Generator  ↵
                        Capture filters
   Version 3.05         Trigger
                        Capture            ↵
  (C) Copyright         Display            ↵
   1986 - 1991          Files
                          More↓
              Test the cable for short or opens.

      Use the arrow keys to move, or ENTER to do this function

1                                                         10 New
 Help                                                     capture
```

Figure 3–1. Cable test in the main menu of the Ethernet analyzer.

The cable tester keeps testing until you terminate it by pressing Esc. While the test is running, the analyzer does not perform any of its other functions.

As long as the tester is active, the Sniffer analyzer repeatedly updates the display transmitted to the console so that the display shows the cable's current status. As long as it detects no fault, it displays the message `No cable fault found`.

The analyzer can detect a cable fault located between the "monitor" adapter and the transceiver that connects it to the network. It can also detect an open line or a short in the network cabling beyond the transceiver (Figure 3–2 and Figure 3–3). The Sniffer analyzer cannot test for faults, open lines, or shorts on cable segments separated by a bridge or repeater from the segment on which it is located.

Automatic Test at First Capture

Under some circumstances, when you ask to start capture, the analyzer first runs a brief cable test. It does this automatically, without being asked. It runs the automatic test only when this is the first live capture since the analyzer program started.[1]

If the automatic test discovers no fault, there is no display and the analyzer proceeds directly with capture.

[1] If the analyzer was already running when you connected to the server, the analyzer may have been running for a considerable time—even days or weeks.

If the automatic test detects a cable fault, the analyzer reports it, and gives you the choice to proceed with capture or to halt.

Inferring Distance from Time

To locate a fault, the Sniffer analyzer notes the interval between its test pulse and the echo that the fault generates. It reports this time in nanoseconds with a resolution of about ±50 nanoseconds. To estimate the probable distance, it divides the time by an estimate of the speed at which the signal propagates through the cable. The actual speed varies not only with the type of cable but also several other factors, including the type of transceivers and connectors, as well as variations in the network's layout.

The analyzer computes two estimates of the distance to the fault. It bases one estimate on the average propagation speed through standard Ethernet cable, which it assumes is about 0.79*c*. (The constant *c* represents the speed of light in a vacuum, about 0.3 meters per nanosecond.) It bases the other on the average propagation speed through thin cable ("Cheapernet"), which it assumes to be about 0.66*c*. These calculations have no way to allow for variations introduced by the actual conditions on the specific network segment the server is monitoring.

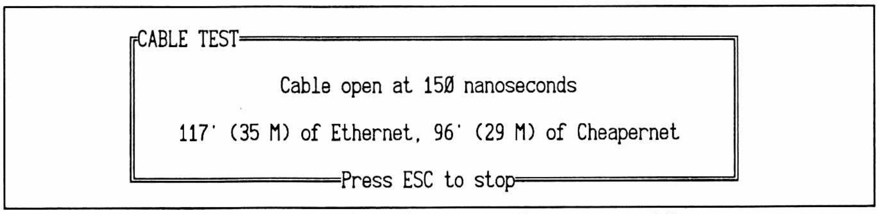

```
CABLE TEST
                 Cable open at 150 nanoseconds

      117' (35 M) of Ethernet, 96' (29 M) of Cheapernet
                       Press ESC to stop
```

Figure 3–2. The Ethernet analyzer's report of an open cable.

Since the resolution of the device producing the timings is to the nearest multiple of 50 nanoseconds, the estimate of distance would have an inherent uncertainty of ±10–12 meters, even if the propagation speed for the segment involved were known precisely. In practice, there is uncertainty about the average propagation speed for the network as a whole and uncertainty about the specific part generating the signal. Also, it may be difficult to know the actual lengths of the cables when they're concealed behind furniture, in the walls, above the ceiling, and so on. Nevertheless, even an uncertain estimate of distance may serve to bracket a range in which the trouble probably lies, and to rule out portions of the network at grossly different distances.

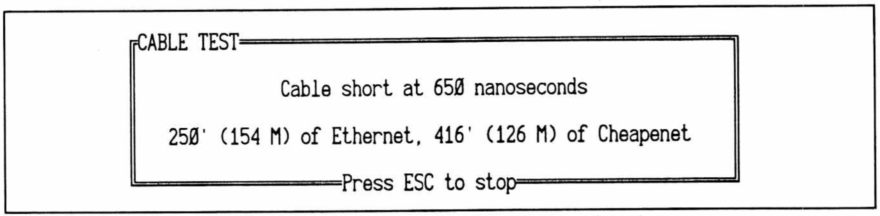

```
CABLE TEST
                    Cable short at 65Ø nanoseconds

  25Ø' (154 M) of Ethernet, 416' (126 M) of Cheapenet
                         Press ESC to stop
```

Figure 3–3. The Ethernet analysis server's report of a short.

To help you interpret the Sniffer analyzer's reports when you really need them, it may be useful to run some calibration tests in advance. Deliberately introduce faults at various known locations. For example, disconnect a cable and leave it without a terminator while you run the cable tester. Note the time and distance that the Sniffer analyzer reports. Note how the report varies during other normal changes (for example, adding or removing a station). It may prove useful to save a table of your observations so you can compare them to the reports you get when the Sniffer analyzer finds a real fault.

Limitations of the Cable Tester

- The Sniffer analyzer readily detects an open line and produces a steady diagnostic display. The reflection characteristic of a short circuit between the center conductor and ground is harder to discriminate from a normal signal, so the server sometimes misses it.
- Certain intermittent defects can produce a jittering display. As it continuously updates the display, the Sniffer analyzer may assign different diagnoses to a continuously changing situation.
- When a collision anywhere on the network coincides with a test pulse, the resulting signal is similar to the pattern produced by an open line. However, a collision produces no more than a momentary flicker.
- Transceivers vary in the way they transmit the test pulse and its echo. This variation in turn produces variation in the behavior of the cable tester. Most transceivers produce useful results, but some do not.
- The procedure for converting time estimates to distance is subject to numerous unpredictable sources of variation.

CHAPTER FOUR: GENERATING TRAFFIC TO LOAD THE NETWORK 4

Chapter 4. Generating Traffic to Load the Network

Chapter Overview

The Traffic Generator is part of the Sniffer analyzer for Ethernet and token ring but not of the analyzer for a WAN/Synchronous link.

The traffic generator permits you to load the portion of the network that the server is monitoring with background traffic. You might do this to observe how other stations respond to delays introduced by a large volume of unrelated traffic, or to test the response of an individual station to heavy traffic of a particular type, or to test the action of gateways and bridges.

Transmissions generated by the traffic generator always obey the network's normal rules for transmitting. For a CSMA/CD network such as Ethernet, that means using the standard collision-detection and back-off algorithms. For a token-passing network, it means waiting for a free token and following the appropriate low-level transmission protocol.

The traffic generator sends the same frame repeatedly (except that it inserts a counter in the frame's data field). You can specify:

- Total length of the frame
- Interval between frames
- Destination address
- Content of the first thirty-two data bytes.

By specifying the first 32 bytes appropriately, you can generate frames that appear to the recipient to have a particular SAP or Ethertype, or frames to which no station should respond.

Don't Swamp the Transport Link Between Console and Server!

If the analysis server's two interface cards are both connected to the *same* network —that is, your console is exchanging control information with the server by way of the network that the server is monitoring— watch out! Flooding the network with traffic may swamp the connection back to the console, making it difficult or even impossible to receive timely reports from the server, and —worse yet— making it difficult or impossible to give the server new instructions. If console and server fail to maintain contact within their time-out parameters, you may lose your connection to the server.

Preparing to Generate Traffic

Figure 4–1 shows the main menu panel from which you select the traffic generator. Pressing Enter at any of the options with ↵ opens a dialog box in which you supply a value or select a destination. Before you start generating, you may set the destination, size, delay, number of frames to be sent, and the content of the data field.

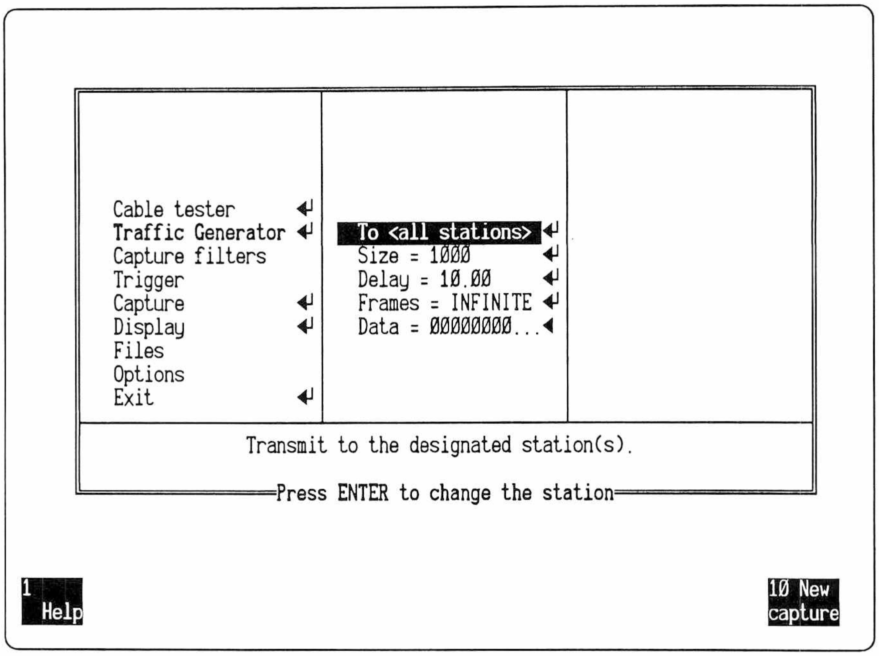

```
Cable tester      ↵
Traffic Generator ↵   To <all stations> ↵
Capture filters       Size = 1ØØØ       ↵
Trigger               Delay = 1Ø.ØØ     ↵
Capture           ↵   Frames = INFINITE ↵
Display           ↵   Data = ØØØØØØØØ...◀
Files
Options
Exit              ↵

            Transmit to the designated station(s).
══════════Press ENTER to change the station══════════

1                                              1Ø New
 Help                                          capture
```

Figure 4–1. The traffic generator menu.

Destination

The traffic generator menu starts with the generated frames' destination address. In the center of Figure 4–1 you can see that it says **To <all stations>**, which is the default. The destination must be a DLC-level address.

To set the destination of the generated frames

1. In the **Traffic generator** menu, move the highlight to **To ...** and press Enter

   Result: The analyzer opens a dialog box showing you a list of DLC addresses currently in the working name table.

2. Move the highlight to the address you want, and press Enter.

   Result: The destination field now shows **To** followed by the name or address you selected.

For example, if you select an address whose name is James, the menu item now says "To James." If you select an address that has no symbolic equivalent (for example, IBM 002FEB), the menu item then says "To IBM 002FEB."

3. To add a new address to the list, move the highlight to **<New station>** and press Enter.

   Result: The analyzer opens a dialog box in which you can write a new DLC address and a symbolic name for it. To select the new address as the destination, return to step 2 (above).

Group destination

You may choose a broadcast address rather than the address of a specific station. Depending on the network, there may also be various classes of group address.

Size of the Generated Frame

The size option sets the total length of the frame to be generated, in bytes. The minimum and maximum permissible lengths depend upon the network, as shown in Figure 4–2.

	Ethernet	Token ring	
		4 Mbps	16 Mbps
Min (bytes)	12	18	18
Max (bytes)	1514	4458	17954

Figure 4–2. Size of frames the traffic generator can send.

On Ethernet, the length of the smallest valid frame is 60 bytes. Shorter frames normally occur only as collision fragments. The Sniffer analyzer is capable of generating short frames. When you specify a length less than 60 bytes, you get a warning that other stations may sense these as collision fragments and therefore may report them as network errors.

To set the size of the frames to be generated

1. In the **Traffic Generator** menu, move the highlight to **Size=** and press Enter.
2. In the dialog box, write the length of the frame and press Enter.

The length is written in decimal, in the same way that the Sniffer analyzer reports the lengths of captured frames: that is, stated as the total length but ignoring the physical header and trailer.

Delay Between Generated Frames

The **Delay** option sets the amount of time the analyzer must wait after it finishes sending one frame until it can start to send another. The interval is the *minimum* interval between the transmitted frames. The actual delay may be longer, since the analyzer may have to wait its turn if other stations are transmitting.

To set the delay between generated frames

1. In the **Traffic Generator** menu, move the highlight to **Delay** and press Enter.
2. In the dialog box, write the length of the delay (in milliseconds) and press Enter.

The dialog box shows you the minimum and maximum values you can enter for the network the server is monitoring, as shown in Figure 4–3.

	Ethernet	Token ring
Minimum delay, milliseconds	0.04	1.0
Maximum delay, milliseconds	1000	1000

Figure 4–3. Minimum interval between consecutive frames sent by the traffic generator.

Number of Frames to Generate

The **Frames** option sets the maximum number of frames to be sent. (You can still tell the analyzer to stop sending before it reaches that number.)

To set the number of frames to be generated

1. In the **Traffic Generator** menu, move the highlight to **Frames** and press Enter.
2. In the dialog box, write the number of frames and press Enter.

You can set any number of frames between 1 and 999999999, or write 0 for infinite, which means that the server keeps transmitting frames until you press Esc to stop it. That's the default.

The Generated Frame's Data Field

The generated frame's destination and source fields occupy the first twelve bytes. All the rest of the frame is considered "data." You can specify what goes into the first 32 of the "data" bytes. The analyzer inserts a counter in the last four bytes of the data field. Thus each generated frame is made up of fields as summarized in Figure 4–4.

Total	**The number of bytes you specify**
Control	Token ring only: two bytes of media access and frame control
Source	Six bytes, generated automatically by the server to identify itself.
Destination	As you select from the list of destinations.
Count	The last four bytes, generated automatically by the analyzer (1 for the first transmitted frame of a series, increased by 1 for each thereafter).
Data	What's in between. The number of available data bytes is therefore Total – (6 + 6 + 4 + (2 if token ring)). Of these, you can specify the first 32 bytes. Any that you don't specify are padded with 00 hex.

Figure 4–4. Fields of a generated frame.

To specify the contents of the generated frames' data field

1. In the **Traffic Generator** menu, move the highlight to **Data =** and press Enter.

   Result: The analyzer opens a dialog box in which you may enter up to 32 bytes of data, in hexadecimal.

2. Type the data you want. When the dialog box opens, the field initially contains 64 zeros (that is, 32 repetitions of 00 hex). Any positions you don't specify remain as 0. Press Enter when your entry is complete.

   Result: The analyzer uses the bytes you specify to fill the first 32 positions of the generated frame's data field.

   If the total length of the data field is less than 32 bytes, the analyzer takes the number needed and ignores the rest. If the total length of the data field is greater than 32 bytes, the analyzer fills the additional space with 00 hex.

   The analyzer inserts the generated frame's sequence number in the last four bytes of the data field. Thus, when the total length of the data field is less then 34 bytes, the sequence number overwrites those positions of the data field.

If your generated frames are sent to a real station and you expect it to read them, it's your responsibility to supply reasonable values in the first part of the data field. This is where an Ethernet recipient will expect to see an Ethertype, and an 802.3 recipient will expect to see length information and an 802.2 header. These are directly affected by the frame's use of routing information.

Effect of Routing Information on the Data Field's Layout

When a frame contains routing information (RI), the RI field starts in the third byte after the source address. But if there is no routing information, that's the start of the 802.2 header or the Ethernet data.

You can't specify what goes in your first 32 bytes until you decide whether your generated frame contains routing information. Moreover, if it does contain routing information, RI is a field of variable length. You must declare its length correctly or your recipient won't be able to locate the fields that are supposed to follow.

<table>
<tr><th colspan="2">Ethernet</th><th>Token ring</th></tr>
<tr><td colspan="2">The hex characters you write specify the first 32 data bytes; that is, those that follow the 6-byte destination and 6-byte source address.</td><td>The hex characters you write specify the first 32 data bytes; that is, those that follow the 1-byte access control, 1-byte frame control, 6-byte destination and 6-byte source address.</td></tr>
<tr><td colspan="2">If you checked the RI option, the variable length source routing information follows the first 2 bytes of user-settable data.</td><td>If you checked the RI option, the first of your 32 bytes contain the source routing information.</td></tr>
<tr><td colspan="2">Interpretation of the first 2 data bytes depends on whether you generate Ethertype or IEEE 802.3 frames (see Figure 4–6).</td><td rowspan="3">The first data byte (following the RI information, if any) identifies the destination SAP, and the second the source SAP. The next one or two bytes are control bytes indicating the type of transmission.</td></tr>
<tr><td><b>Ethertype</b></td><td><b>802.3</b></td></tr>
<tr><td>The first 2 data bytes are the Ethertype. For example, 0800 identifies the IP Ethertype, while 0600 identifies XNS.</td><td>The first 2 data bytes are the 802.2 length, followed by the variable-length RI field, followed by the 802.2 header.</td></tr>
</table>

Figure 4–5. Location of the specifiable data in a generated frame.

The placement of the RI field within an Ethertype frame and within an IEEE 802.3 frame is summarized in Figure 4–6.

	Based on To <xxxxxx>	Automatic	Controlled by what you specify for first 32 bytes			
Ethertype	Destination 6 bytes	Source 6 bytes	Etype 2 bytes	RI 0–32 bytes	Data Remainder	
802.3	Destination 6 bytes	Source 6 bytes	Length 2 bytes	RI 0–32 bytes	802.2 header 3–4 bytes	Data Remainder

Figure 4–6. Position of RI field in Ethertype and IEEE 802.3 frames on Ethernet.

RI Bit in the Source Address of the Generated Frames

A frame originating on a token ring network may include *source routing information*, abbreviated as RI. The RI fields contain a record of each intermediate station that has forwarded the frame. The RI fields are located following the usual DLC source and destination fields. (For a more detailed description of the routing field, see "LAN Option to Interpret or Ignore the RI Bit in a Source Address" on page 2–6.)

To indicate that these optional fields are present, the source address must be modified by forcing a 1 in the bit that (in a destination address) would indicate "broadcast." You can't control the source address of a generated frame. The generated fragment automatically has the source address of the Sniffer server that sent it. However, you can force the analyzer to insert the "RI present" bit.

To turn on the RI bit in the source address of a generated frame

1. In the **Traffic Generator** menu, move the highlight to **Data**.
2. Move rightward to **RI present**. Press Spacebar to toggle X (inactive) to ✓ (active).

If you elect to generate frames with the RI bit turned on, it's your responsibility to include a consistent RI header at the beginning of the data you specify.

Specifying the Length of the RI Field

When the data you provide represent an RI field, you must specify the number of bytes the RI field occupies. The RI field (when present) starts with a two-byte header, followed by from zero to eight 2-byte segment addresses.[1] Thus the total length of the RI field may range from 2 to 18 bytes. The length (mod 32) is encoded in the low-order five bits of the first byte.

[1] These aren't DLC addresses, but segment identifiers adopted by mutual agreement.

Running the Traffic Generator

After you've specified the various parameters—

- Destination
- Size
- Delay
- Quantity
- Data
- RI

—you're ready to tell the analyzer to start generating.

To start the traffic generator

Move the highlight to **Traffic generator** and press Enter.

Result: The analyzer starts transmitting frames as directed. To report its progress, it displays a screen like that shown in Figure 4–7. It updates a counter showing the current frame number, and maintains two thermometer-style bar graphs of frames per second and Kbytes per second transmitted.

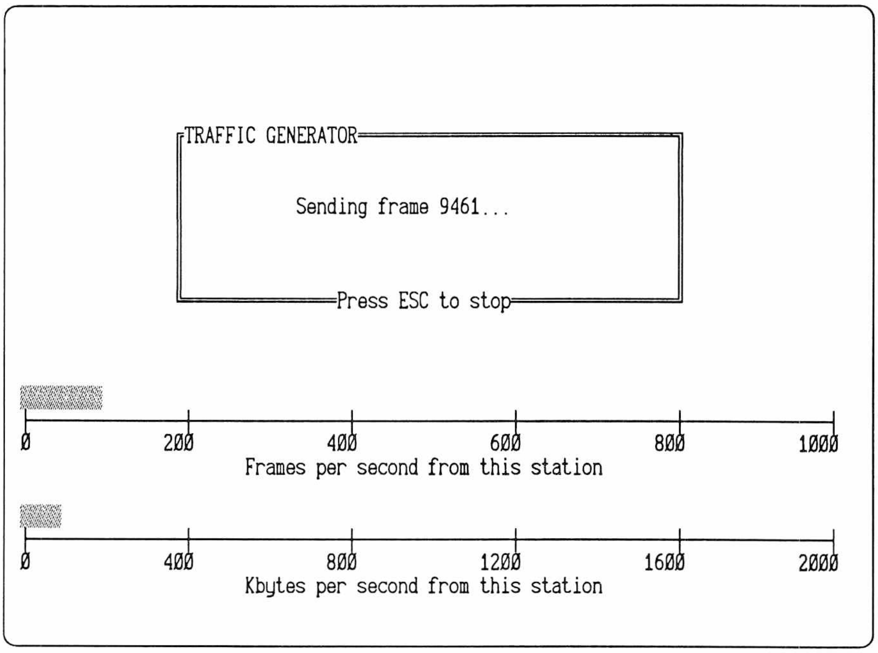

```
TRAFFIC GENERATOR

        Sending frame 9461...

          Press ESC to stop
```

0
200
400
600
800
1000
Frames per second from this station
0
400
800
1200
1600
2000
Kbytes per second from this station

Figure 4–7. Screen displayed at the console while a Sniffer analyzer is generating frames on token ring at 4 Mbps.

While the Sniffer analyzer is generating traffic, it can't perform any of its other functions as an analyzer. However, on a token ring network, it continues to forward incoming frames from its upstream neighbor.

Appearance of the Transmitted Frames

The format depends, of course, on the network. Each LAN frame starts with a header and DLC addresses appropriate to the network. Figure 4–8 shows the token ring frame whose transmission was noted in Figure 4–7, after it has been captured by another Sniffer analyzer. (On a different network, the frame would differ slightly, but would have the same general appearance.)

The first two bytes are the AC and FC bytes of the standard DLC header, visible in the hex view as the characters 18 40.

Six bytes of destination address follow (in this case, C0 00 FF FF FF FF, which means "Broadcast"). Then there are 6 bytes of source address (in this example, 40 00 65 01 00 01, which was the address of the sending Sniffer server).

Then come the 32 bytes you specified for the frame's data field. In the frame shown in Figure 4–8, the data field starts with 00 00 03 00 ... 00 (On token ring, that's the default that the analyzer automatically supplies. It remains in effect if you don't supply your own values for the data field.)

In the detail view, the server interprets the frame. The DSAP and SSAP fields are both 00. Since those do not match any protocol known to the server, it shows the protocol as ???. The server interprets the first 4 bytes as UI, and attaches the explanatory text "Unnumbered information."

In the hex view, the first 3 of those 4 bytes appear highlighted. (The fourth isn't highlighted because the preceding bytes are sufficient to identify the UI command, and the protocol interpreter knows that UI makes no use of the fourth byte.)

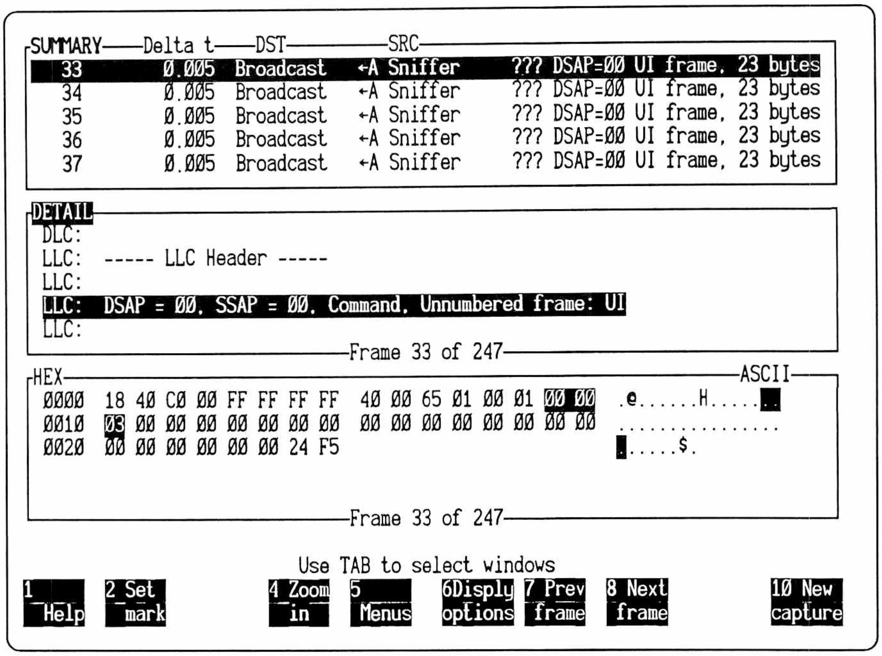

```
SUMMARY——Delta t——DST—————SRC——————————————————————
  33       Ø.ØØ5  Broadcast   ←A Sniffer     ??? DSAP=ØØ UI frame, 23 bytes
  34       Ø.ØØ5  Broadcast   ←A Sniffer     ??? DSAP=ØØ UI frame, 23 bytes
  35       Ø.ØØ5  Broadcast   ←A Sniffer     ??? DSAP=ØØ UI frame, 23 bytes
  36       Ø.ØØ5  Broadcast   ←A Sniffer     ??? DSAP=ØØ UI frame, 23 bytes
  37       Ø.ØØ5  Broadcast   ←A Sniffer     ??? DSAP=ØØ UI frame, 23 bytes

DETAIL
DLC:
LLC:  ----- LLC Header -----
LLC:
LLC:  DSAP = ØØ, SSAP = ØØ, Command, Unnumbered frame: UI
LLC:
                              Frame 33 of 247
HEX———————————————————————————————————————————————————————ASCII———
ØØØØ  18 4Ø CØ ØØ FF FF FF FF  4Ø ØØ 65 Ø1 ØØ Ø1 ØØ ØØ  .e......H......
ØØ1Ø  Ø3 ØØ ØØ ØØ ØØ ØØ ØØ ØØ  ØØ ØØ ØØ ØØ ØØ ØØ ØØ ØØ  ................
ØØ2Ø  ØØ ØØ ØØ ØØ ØØ ØØ 24 F5                           ......$.

                              Frame 33 of 247

                          Use TAB to select windows
```

Figure 4–8. A token ring frame generated by one Sniffer analyzer and captured by another.

Following the standard header and whatever bytes you specify for the data field, the remainder of the frame consists of as many repetitions of 00 (hex) as necessary to fill out the total length you requested. The sequence number occupies the last 4 bytes.

Sequence Numbers

Each frame transmitted by the Sniffer analyzer's traffic generator contains a sequence number. Each time you start the traffic generator, the sequence numbers start at 1. The frame's last four bytes contain the sequence number. If the last four bytes overlap the first 32 data bytes, the sequence number overwrites some of the data you specified for the first 32 bytes.

In Figure 4–8 you can see the frame number. It is encoded as a 32-bit unsigned integer, with the high-order bytes first. Thus, the decimal frame number that appears in Figure 4–7 as 9461 appears in the hex display (Figure 4–8) as 00 00 24 F5.[1]

[1] 24 hex = 36 decimal; F5 hex = 245 decimal.
36 in the 256's position makes 9216.
9216 + 245 = 9461.

CHAPTER FIVE: DISPLAYING AND INTERPRETING CAPTURED FRAMES 5

Network General

Chapter 5. Displaying and Interpreting Captured Frames

Chapter Overview

This chapter describes the various ways you can display and analyze the frames that you've captured. The Sniffer analyzer's most important function —interpreting the various layers of protocol embedded in each frame— is an integral part of the process by which the analyzer displays the contents of the capture buffer.

To select frames for display, the analyzer provides display filters. Their operation is much like the capture filters described in Chapter 2. During display, in addition to the various filters that operate during capture, the analyzer also offers filters based on the frames' interpretation, such as filters for high level protocols or addresses.

The Sniffer analyzer provides three formats for display, called *summary*, *detail*, and *hex*. The chapter describes the three views and their controls, and techniques for browsing through the captured frames.

You can export reports on the selected frames in a variety of formats. One saves an edited trace file, limited to the frames you have selected. Others provide textual reports of the screen displays, or exported files in formats that can be imported by standard spreadsheet or database programs.

This chapter doesn't include the displays that are visible while the server is capturing; they're described in Chapter 2.

Role of the Capture Buffer in Display

For the Sniffer analyzer to display or interpret them, frames must be in the capture buffer. If you've just completed a capture session, the frames you captured are in the buffer. You can start displaying them.

Alternatively, you can load the capture buffer from a file of frames captured earlier. When you load the capture buffer from a file, frames from the file replace those in the buffer. If the buffer now contains frames that you haven't saved and don't want to lose, save them to a file before loading the buffer with a new set. To save the current contents of the capture buffer to a file, or to load a saved file into the capture buffer, start from the **Files** menu; see page 5–69.

To manage how frames are viewed, or to look at particular frames within the capture buffer, start from the **Display** branch of the main menu (see Figure 5–1).

Different Ways to View the Frames After You've Captured Them

The display menu first offers you a choice of four ways to view the captured frames. The options are:

- Summary
- Detail
- Hex
- Two viewports

You can see these options in the right panel of Figure 5–1. The figure shows the screen of an Ethernet Sniffer analyzer, but the display choices are the same for all networks. The options are not mutually exclusive; you can check any combination of them. By default, **Summary** is initially selected and the others are not. Each view is described later in this chapter.

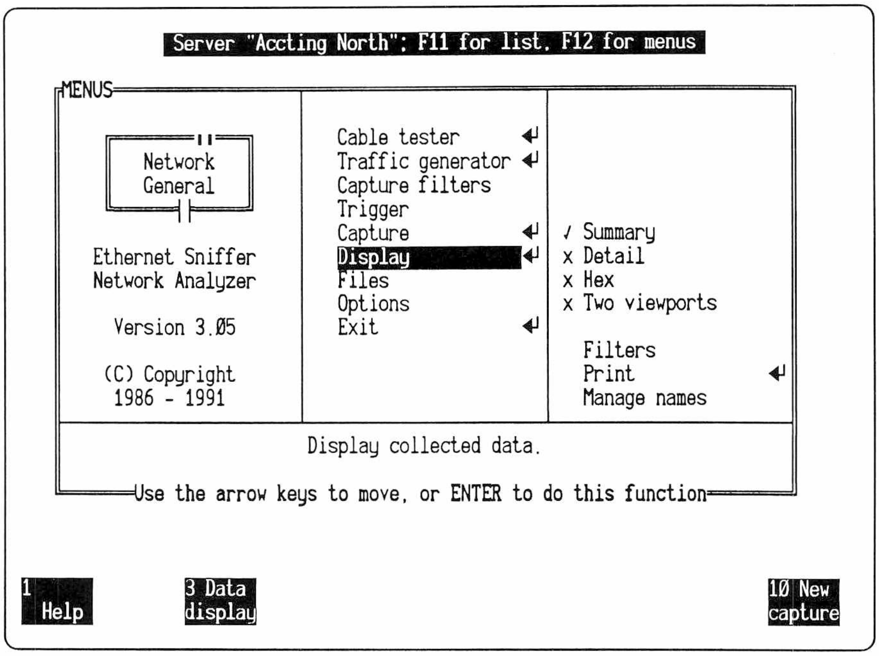

```
            Server "Accting North": F11 for list. F12 for menus

MENUS
  Network              Cable tester       ↵
  General              Traffic generator  ↵
                       Capture filters
                       Trigger
                       Capture            ↵   / Summary
 Ethernet Sniffer      Display            ↵   x Detail
 Network Analyzer      Files                  x Hex
                       Options                x Two viewports
   Version 3.05        Exit               ↵
                                                Filters
  (C) Copyright                                 Print            ↵
   1986 - 1991                                  Manage names

                         Display collected data.
        Use the arrow keys to move, or ENTER to do this function

1               3 Data                                          10 New
 Help           display                                         capture
```

Figure 5–1. The display menu and its branches.

The menu also offers a choice of filters to select frames for display, a print option for reports on selected frames, and an option to manage names to make displays more readable. You can see these choices in the lower portion of the right panel in Figure 5–1.[1]

Once you've selected your display options, you're ready to start display

[1] "Manage names" would ordinarily fall below the lower edge of the panel, and wouldn't become visible until you scroll down.

To display the frames in the capture buffer

1. Press F3.

   (Alternatively, move the highlight back to **Display** and press Enter.)

The sections that follow describe the effects of the various display options.

Setting the Display Filters

Either before you start display, or as display proceeds, you can establish a display filter. Your display filter limits display to a subset of the frames in the capture buffer. Filtering doesn't remove frames from the buffer, but omits them from the display. When some frames are skipped, the frames that remain visible have the same frame numbers as before. Thus, you may see frame 30 followed by frame 35 because the display filter excludes frames 31–34.

When you save the contents of the capture buffer to a file, you may choose to save all frames (regardless of the display filter) or just the frames that the filter accepts.

To set display filters before you start the display itself

1. In the main menu, select Display. Then move the highlight rightward and then vertically to select Filters.
2. Move the highlight vertically to the type of filter you want. Select (in turn, as desired):

   — Address level (page 5–7)

   — Destination class (page 5–8)

   — Station address (page 5–9)

   — Protocol (page 5–13)

   — Pattern Match (page 5–13)

   These are described individually later in this chapter, in sections beginning on the pages noted.

It's easiest to set up a display filter and to select your initial form of display before you actually start displaying frames on the screen. However, it's not hard to alter your filters or your display options after you've started the display.

To change filters or display options once display has started

1. Press F6, **Display options**.

Result: The server opens a temporary panel, superimposed on your current display. Make the changes you want.

2. Press F3 to return to the display (as modified).

You can see the display filters submenu in Figure 5–2. (The figure includes frame quality options that are available only on Ethernet.)

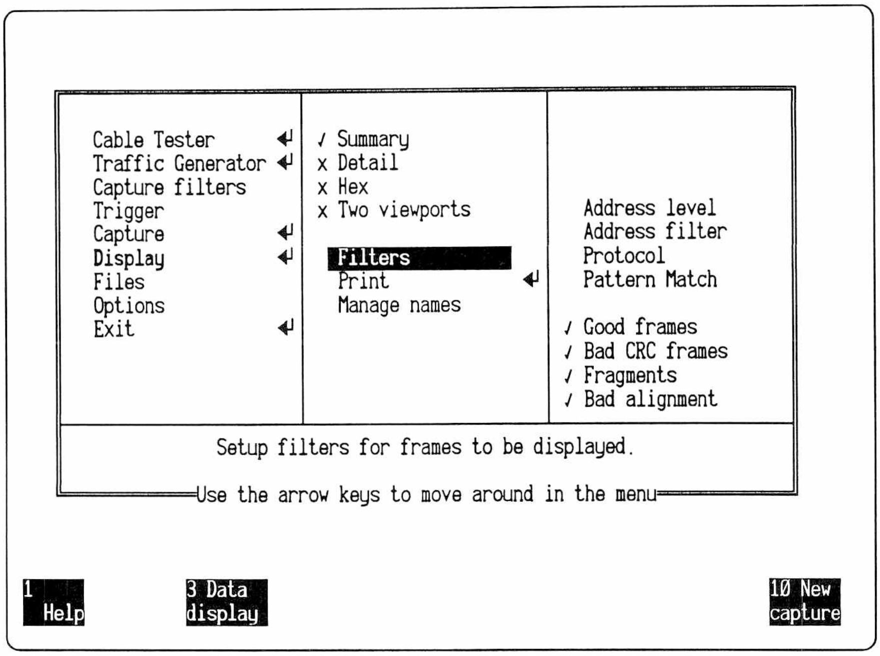

```
Cable Tester       ↵ | ✓ Summary           |
Traffic Generator  ↵ | x Detail            |
Capture filters      | x Hex               |
Trigger              | x Two viewports     |   Address level
Capture            ↵ |                     |   Address filter
Display            ↵ |   Filters           |   Protocol
Files                |   Print           ↵ |   Pattern Match
Options              |   Manage names      |
Exit               ↵ |                     | ✓ Good frames
                     |                     | ✓ Bad CRC frames
                     |                     | ✓ Fragments
                     |                     | ✓ Bad alignment

              Setup filters for frames to be displayed.

          Use the arrow keys to move around in the menu

1        3 Data                                        10 New
 Help    display                                       capture
```

Figure 5–2. The display filters menu.

Criteria for Filtering

The display filter permits the Sniffer analyzer to display a frame that meets all of the following criteria:

Address level — The frame contains an address in one of the indicated protocols.

Destination class — The frame contains a DLC address in the indicated class (such as broadcast or specific).

Station address — The frame contains any of up to four specific addresses at any of the levels selected by the address level filter.

Protocol — The frame contains one or more of the protocols marked with a ✓.

Pattern filter — The frame contains the logical combination of patterns you specified.

Frame defects (On Ethernet, StarLAN and PC Network only.).
The frame exhibits:
- No errors
- Bad CRC
- Fragment
- Bad alignment

—as you place a ✓ beside the categories you want included in the display.

Address Level Filter

Every frame contains both the address of the station from which it just came and the address of the station that is its immediate destination. These addresses are in the frame's lowest level, usually DLC. Frequently a frame contains other addresses as well. For example, it may contain the address of the original source and the address of the ultimate destination. The data field of a lower-level frame thus may include a message written in a higher-level protocol, with its own source and destination, written according to that protocol's rules.

This is almost certainly the case for all frames on a synchronous communications link, and also very likely when frames are being relayed through a gateway between LANs. At the DLC level, a frame's source and destination may be the stations responsible for the current leg of its journey. Within the DLC frame, there may well be addresses in embedded protocols such as XNS, IP, X.400, and so on.

When you set an address level filter, you put a ✓ beside one or more protocols in the panel to the right of address level. The filter accepts only those frames that are addressed in one of the protocols you've checked. To be included in the display, a frame must contain

- A protocol from the indicated set of protocols

  and

- An address in that protocol.

Many protocols require that the frame include an address. But that is not universal. Where the protocol permits both addressed and unaddressed messages, an address level filter accepts a frame only when it actually contains an address.

The Sniffer analyzer's default address level filter has a check at the lowest level of protocol, but at none of the higher levels. Since every frame has a low-level address, the default accepts all frames.

To set an address level filter for display

1. Move the highlight to **Display**, then to **Filters**, then to **Address level.**

2. Move the highlight vertically among the listed protocols. Press Spacebar to toggle between ✓ (include) and X (don't include).

   — Put a check mark beside the protocols you want

   — Remove the check mark from the lowest level, and from other low levels you want to exclude.

Effect on Displayed Names

The **Summary** view shows one address for each frame. When a frame contains multiple levels of address, the **Summary** view shows the highest of the levels you checked in the address level filter. Thus, if you put check marks at the lowest level and at some higher levels, that choice causes the Sniffer analyzer to accept all frames but makes the **Summary** display show higher-level addresses. When the name table contains a symbolic equivalent for that address at that protocol level, the **Summary** view shows the name instead of the numeric address.

Effect on the Name Table

When you edit names (page 5–62), the name table includes an option to add a new station for each of the protocol levels checked in your address level filter.

Destination Class Filter

During display, the destination class filter permits you to include or exclude messages addressed (at any level) to a broadcast destination.[1] You can decide independently whether to include broadcast frames or frames with specific addresses. (Of course, if you exclude them both, you'll have nothing left to look at.)

For each class that you include (broadcast or specific), you also specify the address level (or levels) to be included. The submenu for **Broadcast** and the submenu for **Specific** each contain a list of protocols at which an address may occur (Figure 5–3). By default, initially these are all checked (so that the destination class filter accepts all frames). You turn off the check marks for the levels you don't want. The analyzer accepts a frame if it includes the desired type of address at *any* of the levels checked.

[1] By contrast, during capture the analyzer can filter for broadcast addresses only at the DLC level. Since there are no broadcast addresses at that level on a synchronous link, the WAN/Synchronous analyzer has no capture filter for destination class.

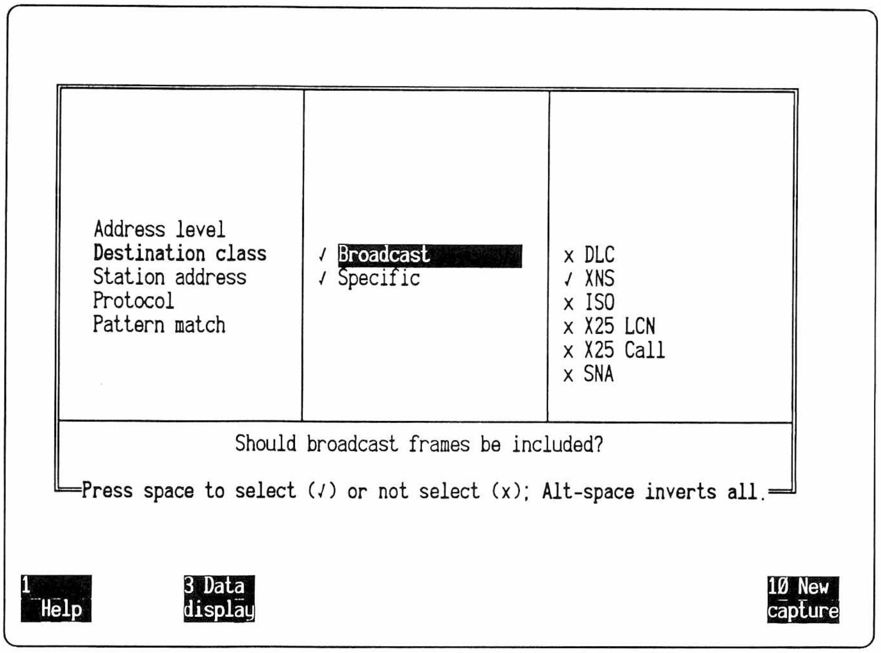

```
Address level
Destination class      ✓ Broadcast      x DLC
Station address        ✓ Specific       ✓ XNS
Protocol                                x ISO
Pattern match                           x X25 LCN
                                        x X25 Call
                                        x SNA

              Should broadcast frames be included?

Press space to select (✓) or not select (x); Alt-space inverts all.

1       3 Data                                   10 New
 Help   display                                  capture
```

Figure 5–3. Address levels within the "Broadcast" filter.

To set a destination class filter for display

1. Move the highlight to **Display**, then to **Filters**, then to **Destination class**.
2. Move the highlight to **Broadcast** or **Specific**. Press Spacebar to toggle between, ✓ (include) and X (don't include).
3. If you checked **Broadcast**, move rightward to the list of address protocols. Use the up or down cursor key to highlight a protocol you want to set. Pres Spacebar to toggle between ✓ (include) and X (don't include), or Alt-Spacebar to reverse the setting for all protocols. (If you leave **Broadcast** without a check mark, it doesn't matter which protocols are checked in Broadcast's submenu.)
4. Similarly, if you checked **Specific**, move rightward to its list of address protocols, and set a check mark at each level that you want to include. (Note that, although the two lists contain the same protocols, your selections of **broadcast** protocols are independent of your selection for **specific** protocols.)

Station Address Filter

A station address filter is formed from some logical combination of up to four pairs of addresses. Address filters for display work just like address filters for capture, but with one important difference. During display, the addresses you specify can be at any of the levels that your protocol interpreters recognize. (By contrast, a capture filter can consider only low-level addresses.)

The procedure for setting an address filter for display is like the procedure for setting one for capture (described starting at page 2–13). However, there are a few additional points because a display filter can include addresses at any level. To review, the procedure is as follows:

To set a station address filter for display

1. In the display menu, select **Filters**. Then select **Station address**.

   Here you may enter from one to four matches, each of which may contain a single address or a pair of addresses.

2. Move the highlight to the first match. Initially, it has the name Match 1.

   You can give a name to each of your four matches. The name has no effect on what the analysis server does. The match names just serve as mnemonic labels.

   To change the match's name, move the highlight to **Match 1** and press Enter; see Figure 2–5, page Figure 2–5.

   Result: The server opens a dialog box in which you can supply a name for this match.

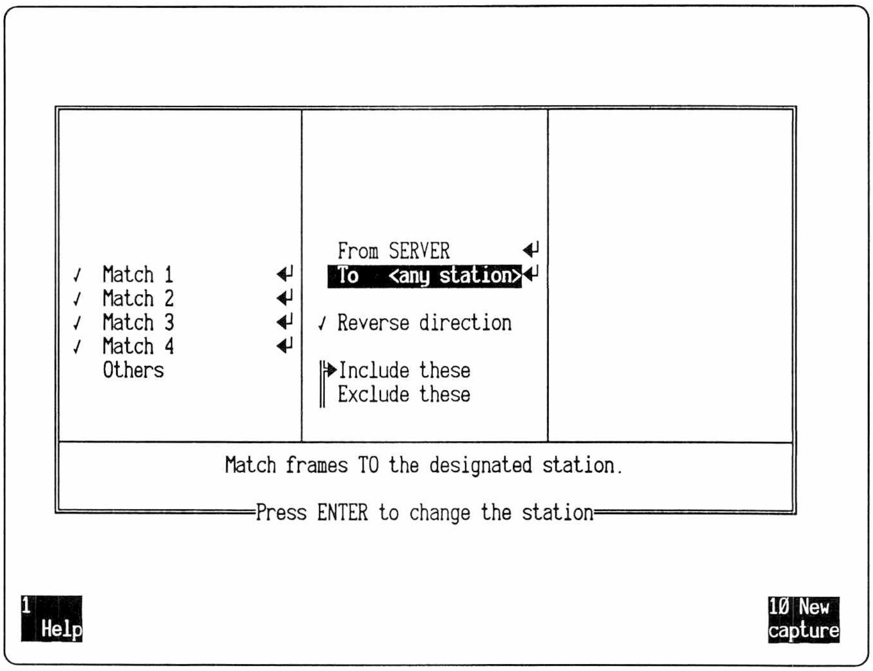

Match frames TO the designated station.

Press ENTER to change the station

Figure 5–4. Menu to select a station address for the filter.

3. To specify the source address for your first match, move the highlight to the entry to the right of Match 1. Initially, it says **From <any station>.**

   Press Enter. The analyzer opens a dialog box showing names and station addresses in the working copy of the name table.

   To select an address for your filter, scroll to the line that contains the address you want and press Enter.

   Result: The analyzer copies the highlighted address into the slot labeled **From**. Instead of saying **From <any station>**, it now says **From <the address you selected>.**

As you scroll through the name table, notice that it contains addresses at various levels (see Figure 5–5). Sometimes the same address may occur at more than one level. When you choose an address, it's important to place the highlight on the address at the level you want.

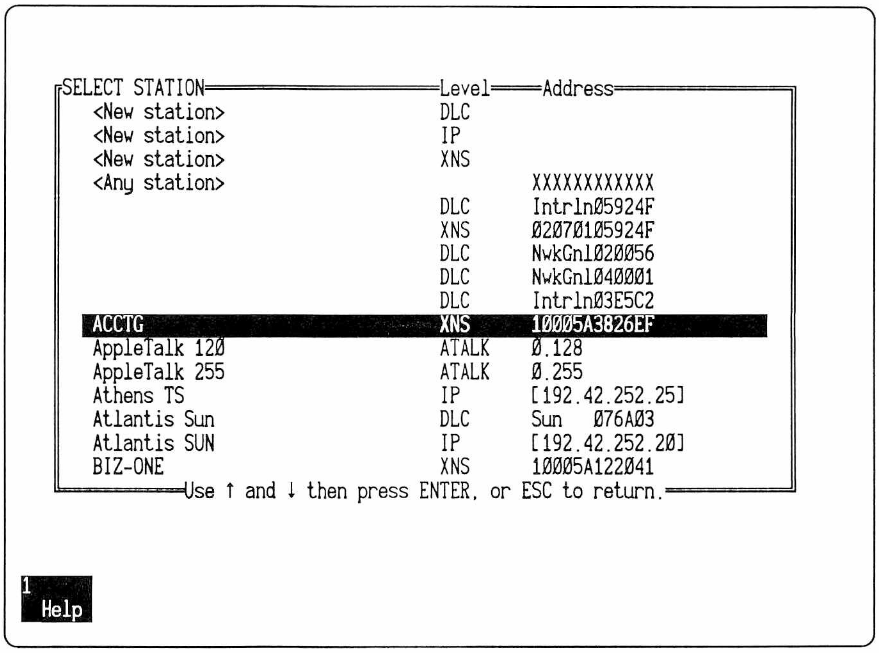

```
SELECT STATION                         Level  Address
   <New station>                       DLC
   <New station>                       IP
   <New station>                       XNS
   <Any station>                              XXXXXXXXXXXX
                                       DLC    Intrln05924F
                                       XNS    02070105924F
                                       DLC    NwkGnl020056
                                       DLC    NwkGnl040001
                                       DLC    Intrln03E5C2
   ACCTG                               XNS    10005A3826EF
   AppleTalk 120                       ATALK  0.128
   AppleTalk 255                       ATALK  0.255
   Athens TS                           IP     [192.42.252.25]
   Atlantis Sun                        DLC    Sun   076A03
   Atlantis SUN                        IP     [192.42.252.20]
   BIZ-ONE                             XNS    10005A122041
           Use ↑ and ↓ then press ENTER, or ESC to return.
```

Figure 5–5. Dialog box in which to select a station address for display filters.

It includes a new address entry for each protocol checked in the address level filter.

4. To specify a destination address, move the highlight to **To** and press Enter. Supply an address in the same way as Step 3.

5. To indicate whether this match also applies to traffic in the reverse direction, move the highlight to **Reverse direction.** Press Spacebar to toggle between √ (include) and x (exclude).

6. To indicate whether frames identified by this match should be included or excluded, move the highlight to the radio control **Include these** or **Exclude these**, and press Spacebar at the one you want.
7. Repeat steps 1 through 6 for up to four matches.
8. To specify what the analyzer should do with frames not covered by your matches, move the highlight to **Others** and then to **Include** or **Exclude**, and there press Spacebar.

Order in Which Matches Are Checked

See "Taking Advantage of the Order in Which Matches are Checked" on page 2–17. Although it's about address filters during capture, it applies equally to address filters during display.

Adding Entries to the Working Name Table

If the address you want isn't yet in the name table, you must first add it to the table and then select it. There are two ways to add new addresses:

- Let the Sniffer analyzer scan the capture buffer and add new addresses automatically
- Add the new addresses manually.

Automatic Scanning for New Addresses

The easiest way to make sure that all addresses in the capture buffer are in the name table is to have the Sniffer analyzer scan the entire contents of the capture buffer. It does that automatically the first time you use display following capture (see page 5–65).

Adding Addresses Manually

To supply a new address for a station address filter

1. Highlight the line that says **<New station>** for the address level you want.

   The name table contains a **<New station>** entry for each of the levels checked in the address level filter. If the name table lacks a **<New station>** entry for the level you want, probably you haven't yet checked that level in the address level filter. Return to the address level filter, adjust it appropriately, and then go back to defining your station address filter.
2. With **<New station>** highlighted, press Enter.

3. The server opens a dialog box in which to write the address and a name for it.
4. Press Enter. The analyzer records the name in its working copy of the name table and at the same time makes it the effective address of the match you are defining.

Unnamed addresses are transitory

The working name table lasts only until you exit the Sniffer analyzer program. You can save the name table by selecting manage names and then save names. But that only saves addresses that have names. Addresses that lack symbolic equivalents are dropped. Be sure you have assigned a name to each address you want to preserve.

Protocol

You can filter for frames that contain a protocol that interests you. Set a check mark beside the name of each protocol that you wish to include in the display. The filter includes a frame in the display when it contains *any* of the protocols you check.

During display, the Sniffer analyzer can filter for any protocol that its protocol interpreters recognize. (However, during capture, the Sniffer analyzer filters only on the lowest-level protocols.)

To select protocols for the display filter

The procedure to mark protocols in the display filter is the same as the procedure to mark protocols in the capture filter (see the discussion starting at page 2–18). However, for display, the list of possible protocols is longer. That's because it includes any of the protocols interpreted by your protocol interpreter suites, and not just protocols at the DLC level.

A quick way to select only one protocol is as follows:

- Move the highlight to the one you want.
- Press Alt-Spacebar to change ✓ to x for all protocols.
- Then press Spacebar to change x to ✓ for the protocol you want.

Pattern Match

You can set a filter to accept frames that contain a set of patterns. The set may be built from a logical combination of up to eight component patterns. Each component pattern is a specific sequence of characters at a specific location (each described as hex, binary, or text).

To set up a pattern match filter for display

The procedure for specifying a set of patterns for the display filter is exactly the same as the procedure for specifying a set of patterns for the capture filter (see the discussion starting at page 2–19). The only difference is that, for display patterns, you highlight **Pattern match** in the **Display filters** menu rather than in the **Capture filters** menu.

Defective Frame

On Ethernet, the **Display Filters** menu includes four additional options. (They also appear in the **Capture Filters** menu.) The four are:

- Good frames (that is, frames having none of the following defects).
- Bad CRC frames.
- Fragments.
- Misaligned Frames.

On a WAN, the only filterable defect is **Bad CRC**. On token ring, the "monitor" interface card does not pass defective frames to the Sniffer server.

To set filters for frame defects

1. In the **Display filters** menu, move the highlight downward to the last section.

   Result: You'll see the list of defect categories for the network the analyzer is monitoring. A ✓ mark indicates that frames in the category should be accepted, X indicates that they should be excluded.

2. Move the highlight to each category you wish to change. Press Spacebar to toggle between ✓ and X.

Three Ways to View Frames

There are three ways to view a frame. You can have just one view on the screen or any combination of the three. These views can be focused on a single frame in the capture buffer, or at two different frames. The three views are:

Summary view

A short form of the hexadecimal and **Detail** listings condensed so that each level fits on a single line. You may elect to show all levels of interpretation or only each frame's highest level.

Detail view — The protocol is identified and standard fields within it are labeled and explained. Station addresses are replaced by their symbolic equivalents according to the address table you supplied in the default startup file (see page 6–10ff). Since a low-level frame may contain higher-level frames within it, a single frame may require several levels of interpretation.

Hexadecimal view — All bytes of the entire frame are shown, accompanied by a transliteration to ASCII or EBCDIC characters.

How the three views are positioned

When you have more than one view open, they're always in the same order. **Summary** is always above the others. Hex is always below the others.

Dual Viewports

In addition to these three views, you also have the option to split the screen vertically into two independent viewports. The two sides contain the same combination of **Summary, Detail,** and **Hex**. Having two viewports permits you to scroll independently on the left and right sides of the screen. That way you can keep a frame in view on one side while on the other side you scroll forward or back to look for related frames.

Zoom

To get an enlargement of one of the views, press F4, Zoom in. The active view then occupies the entire area, temporarily concealing the others. To return to the multi-panel overview, press F4 again.

Active Panel

The view that contains the cursor is said to be *active*. The border of the active panel is shown intensified (brighter or in a contrasting color, depending on the monitor). The Tab key moves the highlight from one panel to the next, activating the panel in which it arrives. Shift Tab moves the highlight in the reverse direction.

The Summary View

The **Summary** view gives you a condensed view of your captured frames. Each is reduced to a single line or a few lines. The **Summary** view is the only view that shows several frames at once. Although each frame is abbreviated and condensed, you can see at a glance the sequence and context of the frames. You can then examine individual frames in greater detail or skip over them.

You have several choices on the **Summary** submenu (visible in the right panel of Figure 5–6).

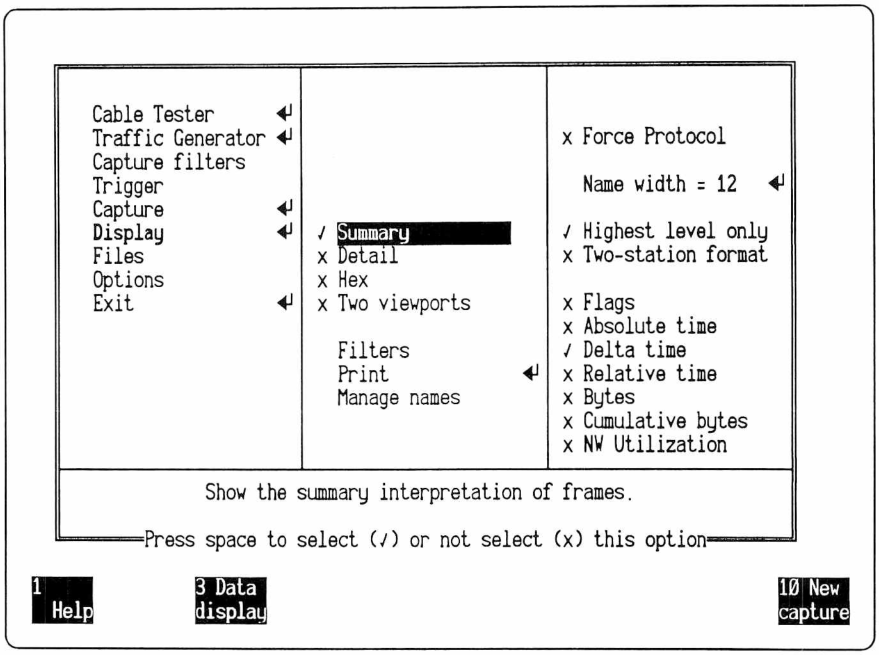

```
Cable Tester      ↵                          x Force Protocol
Traffic Generator ↵
Capture filters                                Name width = 12  ↵
Trigger
Capture           ↵                          ✓ Highest level only
Display           ↵  ✓ Summary               x Two-station format
Files                x Detail
Options              x Hex                   x Flags
Exit              ↵  x Two viewports         x Absolute time
                                             ✓ Delta time
                       Filters               x Relative time
                       Print           ↵     x Bytes
                       Manage names          x Cumulative bytes
                                             x NW Utilization

              Show the summary interpretation of frames.
        Press space to select (✓) or not select (x) this option

1 Help        3 Data display                              10 New capture
```

Figure 5–6. Options for the **Summary** *display.*

Names in the Summary View

In the summary view, when the Sniffer analyzer knows a station's name, it uses it. But when an address has no symbolic equivalent in the name table, the analyzer displays the station's numeric address. In general, a numeric address is shown in the conventional format for its type. For example, a DLC address is shown in hexadecimal, but an IP address is shown as [n.n.n.n], and so on.

On Ethernet or token ring, each station's DLC address contains six bytes. Six bytes can be written as 12 hexadecimal digits. This is the default width of the name field in the **Summary** display for all Sniffer analyzers (regardless of network).

Manufacturer IDs

Where the analysis server shows a 6-byte DLC address, it attempts to interpret the first three bytes as the name of the NIC's manufacturer. When it is able to find the manufacturer's code in its table of manufacturers, it replaces the first six characters of the station address with an ASCII abbreviation of the manufacturer's name.

The file containing names for manufacturer IDs is named STARTUP.*xx*D (where *xx* stands for the two-letter network abbreviation (TR, EN, SY). See "Table of Manufacturer ID Codes and Abbreviations" on page 6–17.

Higher-Level Addresses

To have the Sniffer analyzer use higher-level protocol addresses instead of DLC addresses, you must check the desired protocol levels in the address level filter; see page 5–7.

Width of the Name Field

You may change the width of the name field in the **Summary** display. The name table permits names up to a maximum of 31 characters. If you use long names, you may want to make the name field wider to accommodate them. This may push some other fields off the screen, but you can always scroll sideways to see what's there. If you use an external screen, you may be able to display more rows or columns. The analysis server takes advantage of extra columns if they're available.

The name field's smallest permissible width is six characters. When the display includes a name longer than the name field, the Sniffer analyzer truncates the name and replaces the last two visible characters with dots (to show that the name has been truncated). For example, if you squeeze the 10-character name **FileServer** into an 8-character field, it appears as **FileSe..** (that is, six characters and two dots).

Multiple Levels vs. Highest-Level Only

You can toggle the option **Highest level only**. That is, either it's checked or it's not checked.

Highest level only

✓	The **Summary** view shows only one line for each frame, summarizing the highest-level protocol that the frame contains.
X	The **Summary** view shows a separate line for each protocol the frame contains.[1] The levels are arranged with the outermost (lowest) level on top. The protocols are color-coded by level (see Figure 5–18, page 5–33).

[1] In X Windows, there is a separate line for each protocol of each message within the frame.

To display all levels of protocol in the **Summary** *view*

1. Move the highlight to **Display**, then to **Summary**, then to the panel to its right.
2. Move the highlight to **Highest level only** and press Spacebar to toggle from √ (active) to X (inactive).

Two-Station Format

Analysis often focuses on the interchange between a pair of stations. To make the dialog clearer, use filters to show those stations only and elect **two station** format. Two station format shows transmissions from one station on the left side of the screen, and transmissions from the other station on the right. An example is shown in Figure 5–7. (In this figure, the display is shown with highest level only turned off, so there is a separate line for each level of protocol.)

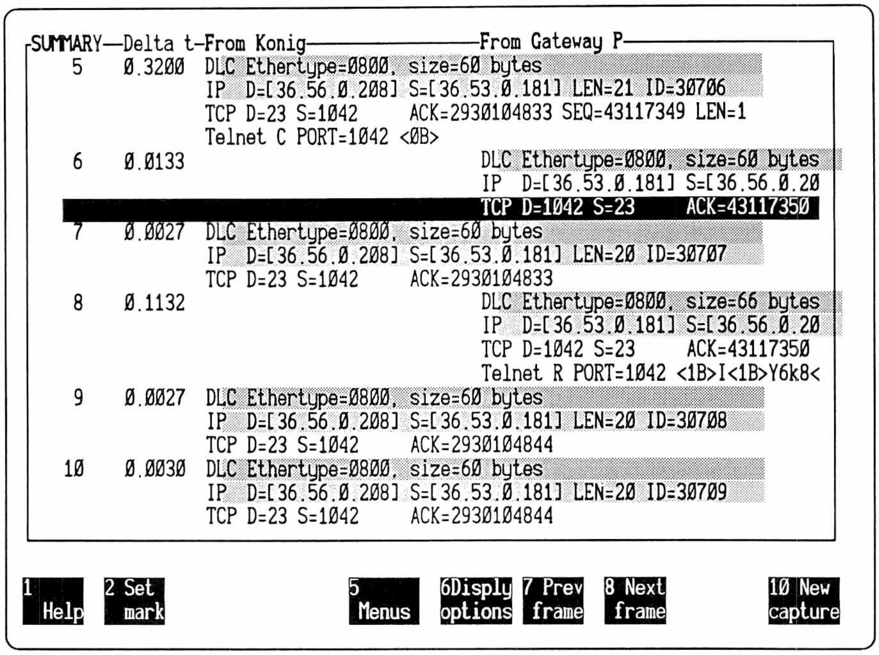

```
SUMMARY—Delta t-From Konig——————————————From Gateway P——————————
    5     0.3200  DLC Ethertype=0800, size=60 bytes
                  IP  D=[36.56.0.208] S=[36.53.0.181] LEN=21 ID=30706
                  TCP D=23 S=1042     ACK=2930104833 SEQ=43117349 LEN=1
                  Telnet C PORT=1042 <0B>
    6     0.0133                                   DLC Ethertype=0800, size=60 bytes
                                                   IP  D=[36.53.0.181] S=[36.56.0.20
                                                   TCP D=1042 S=23      ACK=43117350
    7     0.0027  DLC Ethertype=0800, size=60 bytes
                  IP  D=[36.56.0.208] S=[36.53.0.181] LEN=20 ID=30707
                  TCP D=23 S=1042     ACK=2930104833
    8     0.1132                                   DLC Ethertype=0800, size=66 bytes
                                                   IP  D=[36.53.0.181] S=[36.56.0.20
                                                   TCP D=1042 S=23      ACK=43117350
                                                   Telnet R PORT=1042 <1B>I<1B>Y6k8<
    9     0.0027  DLC Ethertype=0800, size=60 bytes
                  IP  D=[36.56.0.208] S=[36.53.0.181] LEN=20 ID=30708
                  TCP D=23 S=1042     ACK=2930104844
   10     0.0030  DLC Ethertype=0800, size=60 bytes
                  IP  D=[36.56.0.208] S=[36.53.0.181] LEN=20 ID=30709
                  TCP D=23 S=1042     ACK=2930104844
```

Figure 5–7. Two station format for the **Summary** *view.*

In two station format, the source and destination fields are omitted. Instead, there are two columns, headed **From** *xxx* and **From** *yyy*. A frame from the station on the left is assumed be addressed to the station on the right, and vice versa.

To switch to Two-station format

1. Move the highlight to **Display**, then to **Summary**, then to the panel to the right.
2. Move the highlight to **Two-Station Format**. Press Spacebar to toggle between X (inactive) and √ (active).

3. Useful but not essential: set filters so that only the two stations of interest appear in the display.

Which Stations Appear in Two-Station Format

The source that appears earliest in the capture buffer is shown on the left; the source that appears second appears on the right.

To adjust the separation between the columns, change the width of the name field. To adjust the number of columns devoted to time displays, change your selection of time measurements (see Figure 5–8). To adjust the horizontal position of the display on screen, use the horizontal cursor position keys to scroll sideways.

Frames that Don't Fit the Two-Station Paradigm

If your display filter accepts frames from other sources or with other destinations, the analysis server displays the additional frames in the usual format. That is, for those frames only, the display uses the full width, and shows source and destination explicitly. Since this is inconsistent with the two-station format, it's usually preferable to combine two-station format with a filter that excludes other traffic.

Options to Show Time, Network Utilization, and Frame Size

For every frame, the **Summary** display permits you to include or omit various indicators of the time or the flow of data. The alternatives are listed in Figure 5–8 and defined in Figure 5–10.

Option	Effect	Whether Required	Columns used (see note)
Flags	Widen **flags** field to 6 characters, to allow for flags in addition to T and M.	One col. required	No flags: 1 Flags: 6
Frame number	Show unique serial number for each frame in the buffer.	Required	5 (+1)
Absolute time	Show time at which the frame was received.	Optional	14 (+1)
Delta time	Show interval since the preceding displayed frame arrived.	Optional	9 (+1)
Relative time	Show interval since the marked frame.	Optional	9 (+1)
Bytes	Show length of this frame.	Optional	5 (+1)
Cumulative bytes	Show length of this frame and of all frames between it and marked frame.	Optional, but mutually exclusive.	6 (+1)
Network usage	Show bytes in this frame as percent of theoretical capacity in the interval.		7 (+1)
Note: The entry 5(+1) means that the field assigns five columns for text followed by one blank column as a separator.			

Figure 5–8. Displays at the left side of the Summary *view.*

Flags Option

The leftmost column of the **Summary** display always contains the T or M flags. When you select the flags option, the field is widened from one column to six. Flags appear as letters of the alphabet, left-justified in the field. Flags defined by the Sniffer analyzer are shown in Figure 5–9.

Universal	M	Mark. The reference frame for relative time or cumulative bytes. To set the mark on the current frame, press F2.
	T	Trigger. Marks the frame whose detection prompted the end of capture.
Optional	A	Alignment error (Ethernet). Marks a frame that reached the network interface card with a number of bits not divisible by 8. Such a frame may arise when the sender detects a collision and aborts transmission. The NIC discards the trailing bits before passing the frame to the Sniffer analyzer.
	A	Abort (WAN/Synchronous). Preceding the frame thus marked, one or more frames were aborted by the sender
	C	CRC (error cyclic redundancy check). Marks a frame whose CRC does not agree with the actual bytes received, suggesting that it contains invalid characters.
	F	Fragment (Ethernet): Frame less than minimum legal length.
	L	Lost frame (Ethernet). Preceding the frame thus marked, one or more frames reached the network interface card but were lost before being passed to the Sniffer analyzer.
	O	Overrun (Ethernet). Preceding the frame thus marked, data was lost because of an overrun during transfer from NIC to CPU by direct memory access.

Figure 5–9. Codes used in the flags field of the **Summary** *display.*

Measures of Time, Network Utilization, and Frame Size

To examine a network's throughput, you need data on the volume and timing of transmissions. In its **Summary** display, the Sniffer analyzer provides several alternative measures of time and throughput, selectable to match the situation. You can include various combinations of these measures in the same display (see Figure 5–11). Additional measures require additional columns. This may push other parts of the display out of the visible portion of the panel, but you can always scroll to see the portion thus concealed. The measures are listed in Figure 5–10.

Absolute time	When the Sniffer analyzer completes reception of a frame, it attaches a timestamp. The timestamp records the time according to the Sniffer server's internal clock at the moment the Sniffer analyzer recorded the end of the frame. All displays of time are computed from the absolute value recorded with each frame. Absolute time is displayed as hours, minutes, and seconds to the nearest millisecond (on the token ring), or to the nearest tenth of a millisecond (on networks other than token ring). That's also how time is shown in the **Detail** display.
Delta time	The time shown is the interval between the current frame's timestamp and the timestamp of the preceding frame in the display. Delta time is shown with the same precision as absolute time. Note that because it's the interval to the preceding displayed frame, frames that are not displayed don't affect delta time.
Relative time	The time shown is the difference between the current frame's timestamp and the timestamp of the reference frame. Relative time is shown with the same precision as absolute time. The reference frame is marked by the letter M to the left of the frame number. When you first display the buffer, the first frame is marked. You can mark a frame (thereby removing the mark from any other frame) by pressing F2 while you have the frame highlighted. Once you've marked a reference frame, you can find it quickly by the display option **Jump to mark** (see 5–39).
Bytes	The number shown is the total number of bytes in the frame, not including the CRC frame.
Cumulative Bytes	The number shown is the sum of the lengths of the displayed frames, from the reference frame through the current frame (including both). When you haven't marked a reference frame, the Sniffer analyzer counts from the first displayed frame.
Network Utilization	The number shown is an estimate of the percentage of the network's bandwidth devoted to transmitting the displayed frame (and perhaps those preceding and following it). Basically, the measure is $\frac{\text{100 × bytes in all frames accepted during the interval}}{\text{Theoretical maximum that could be transmitted during the interval}}$ The interval is a time window centered around the frame. You can set the size of the interval to 1, 10, 100 or 1000 milliseconds. For example, if you pick 100 millisecond intervals, the utilization for a frame that arrived at 13:27:06.100 is based on the number of bytes in frames whose arrival times ranged from 13:27:06.050 to 13:27:06.150. Utilization is a moving average. With a small interval, you'll see larger momentary fluctuations; a larger interval smooths them out. Any measure of network utilization must be based on a time window, whether described explicitly or not. Viewed without window averaging, a network is always either 100% busy (when a frame is being transmitted) or 0% busy (when no frame is being transmitted).

Figure 5–10. Descriptions of the optional displays for time, traffic, and density.

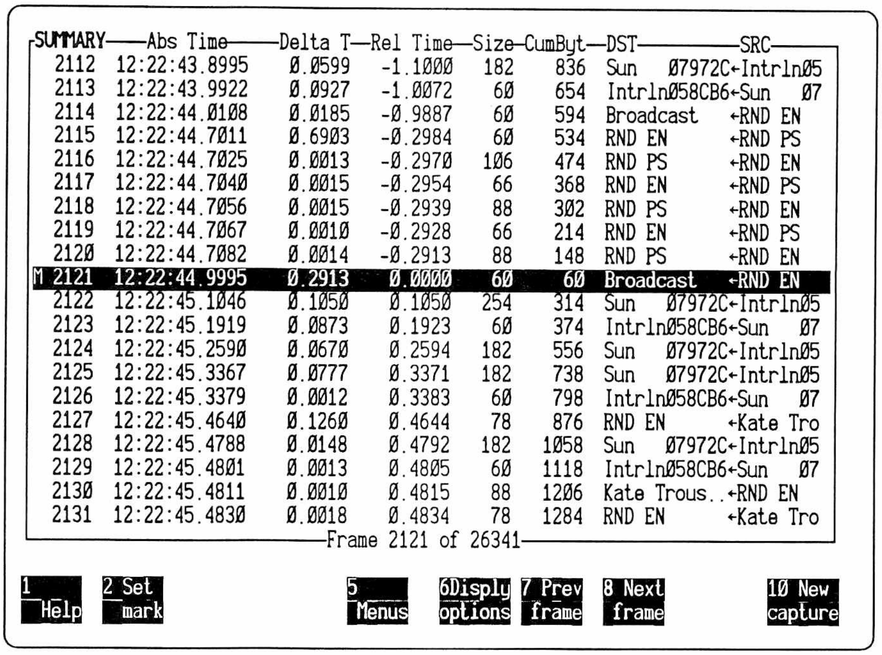

```
SUMMARY——Abs Time———Delta T—Rel Time—Size—CumByt—DST——————SRC———
  2112  12:22:43.8995    Ø.Ø599   -1.1ØØØ  182    836  Sun   Ø7972C←IntrlnØ5
  2113  12:22:43.9922    Ø.Ø927   -1.ØØ72   6Ø    654  IntrlnØ58CB6←Sun   Ø7
  2114  12:22:44.Ø1Ø8    Ø.Ø185   -Ø.9887   6Ø    594  Broadcast    ←RND EN
  2115  12:22:44.7Ø11    Ø.69Ø3   -Ø.2984   6Ø    534  RND EN       ←RND PS
  2116  12:22:44.7Ø25    Ø.ØØ13   -Ø.297Ø  1Ø6    474  RND PS       ←RND EN
  2117  12:22:44.7Ø4Ø    Ø.ØØ15   -Ø.2954   66    368  RND EN       ←RND PS
  2118  12:22:44.7Ø56    Ø.ØØ15   -Ø.2939   88    3Ø2  RND PS       ←RND EN
  2119  12:22:44.7Ø67    Ø.ØØ1Ø   -Ø.2928   66    214  RND EN       ←RND PS
  212Ø  12:22:44.7Ø82    Ø.ØØ14   -Ø.2913   88    148  RND PS       ←RND EN
M 2121  12:22:44.9995    Ø.2913    Ø.ØØØØ   6Ø     6Ø  Broadcast    ←RND EN
  2122  12:22:45.1Ø46    Ø.1Ø5Ø    Ø.1Ø5Ø  254    314  Sun   Ø7972C←IntrlnØ5
  2123  12:22:45.1919    Ø.Ø873    Ø.1923   6Ø    374  IntrlnØ58CB6←Sun   Ø7
  2124  12:22:45.259Ø    Ø.Ø67Ø    Ø.2594  182    556  Sun   Ø7972C←IntrlnØ5
  2125  12:22:45.3367    Ø.Ø777    Ø.3371  182    738  Sun   Ø7972C←IntrlnØ5
  2126  12:22:45.3379    Ø.ØØ12    Ø.3383   6Ø    798  IntrlnØ58CB6←Sun   Ø7
  2127  12:22:45.464Ø    Ø.126Ø    Ø.4644   78    876  RND EN       ←Kate Tro
  2128  12:22:45.4788    Ø.Ø148    Ø.4792  182   1Ø58  Sun   Ø7972C←IntrlnØ5
  2129  12:22:45.48Ø1    Ø.ØØ13    Ø.48Ø5   6Ø   1118  IntrlnØ58CB6←Sun   Ø7
  213Ø  12:22:45.4811    Ø.ØØ1Ø    Ø.4815   88   12Ø6  Kate Trous..←RND EN
  2131  12:22:45.483Ø    Ø.ØØ18    Ø.4834   78   1284  RND EN       ←Kate Tro
———————————————————Frame 2121 of 26341——————————————————————————————
```

Figure 5–11. Display showing several measures of time.

The Detail View

The **Detail** view presents a complete interpretation of a frame, labeling each field and decoding the value of the field's parameters. It also shows certain information not contained within the frame, such as the time on the Sniffer server's clock when the frame arrived.

Figure 5–12 shows the text of a **Detail** display. The **Detail** view frequently requires many lines. On screen, you can scroll to the parts not immediately visible. Figure 5–12 shows the entire **Detail**, as it appears when you print it. The example is a TCP/IP frame transmitted over Ethernet, but the general format is similar for any network. The left margin of the **Detail** view indicates the protocol governing that portion of the interpretation.

Scrolling

In any panel, when the information takes more space than can fit on screen, you can use the cursor keys to scroll hidden text into view.

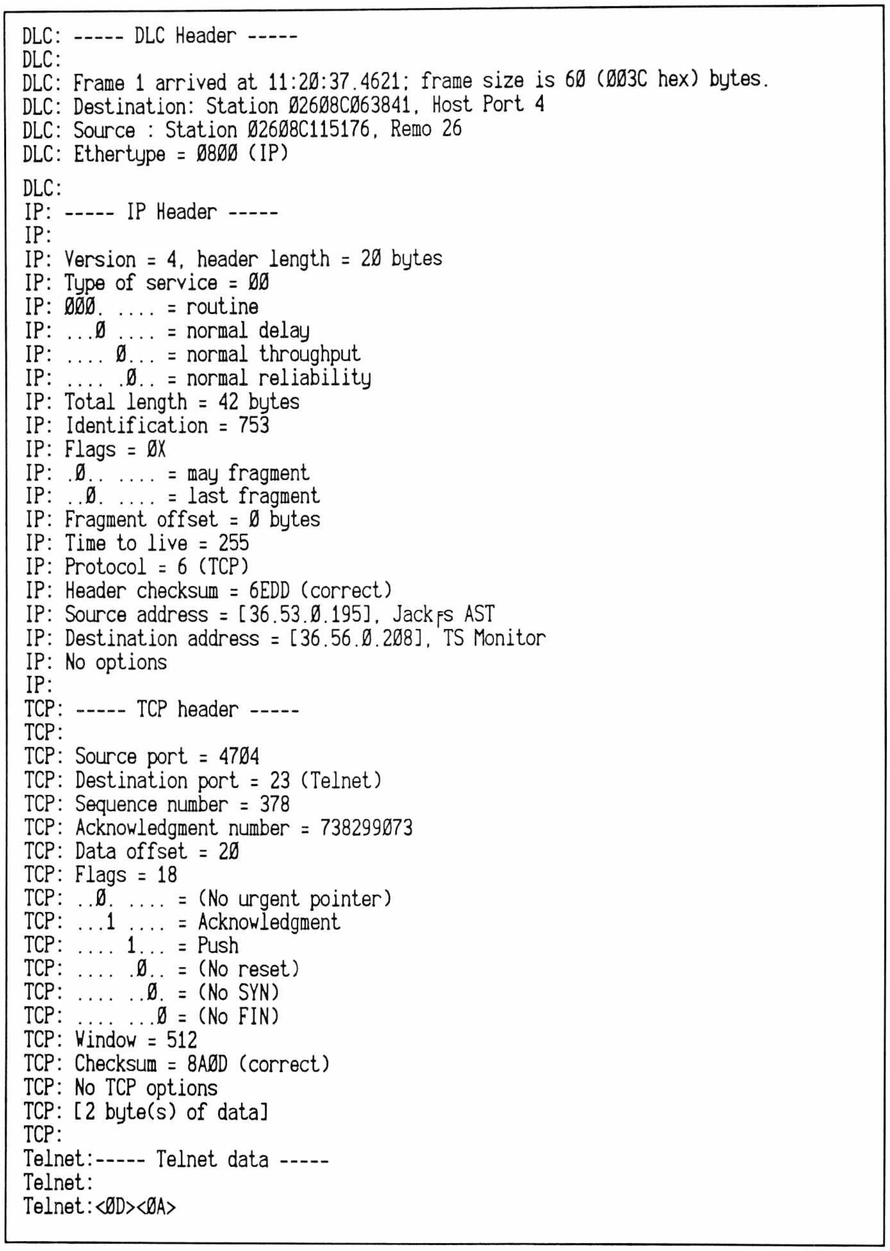

```
DLC: ----- DLC Header -----
DLC:
DLC: Frame 1 arrived at 11:2Ø:37.4621; frame size is 6Ø (ØØ3C hex) bytes.
DLC: Destination: Station Ø26Ø8CØ63841, Host Port 4
DLC: Source : Station Ø26Ø8C115176, Remo 26
DLC: Ethertype = Ø8ØØ (IP)
DLC:
IP: ----- IP Header -----
IP:
IP: Version = 4, header length = 2Ø bytes
IP: Type of service = ØØ
IP: ØØØ. .... = routine
IP: ...Ø .... = normal delay
IP: .... Ø... = normal throughput
IP: .... .Ø.. = normal reliability
IP: Total length = 42 bytes
IP: Identification = 753
IP: Flags = ØX
IP: .Ø.. .... = may fragment
IP: ..Ø. .... = last fragment
IP: Fragment offset = Ø bytes
IP: Time to live = 255
IP: Protocol = 6 (TCP)
IP: Header checksum = 6EDD (correct)
IP: Source address = [36.53.Ø.195], Jack┌s AST
IP: Destination address = [36.56.Ø.2Ø8], TS Monitor
IP: No options
IP:
TCP: ----- TCP header -----
TCP:
TCP: Source port = 47Ø4
TCP: Destination port = 23 (Telnet)
TCP: Sequence number = 378
TCP: Acknowledgment number = 738299Ø73
TCP: Data offset = 2Ø
TCP: Flags = 18
TCP: ..Ø. .... = (No urgent pointer)
TCP: ...1 .... = Acknowledgment
TCP: .... 1... = Push
TCP: .... .Ø.. = (No reset)
TCP: .... ..Ø. = (No SYN)
TCP: .... ...Ø = (No FIN)
TCP: Window = 512
TCP: Checksum = 8AØD (correct)
TCP: No TCP options
TCP: [2 byte(s) of data]
TCP:
Telnet:----- Telnet data -----
Telnet:
Telnet:<ØD><ØA>
```

Figure 5–12. Detail view of a sample frame.

When you send the display to a printer, the SniffMaster console prints the entire text of the view (or views) you've selected, regardless of the boundaries of the screen.[1]

[1] Depending how you've configured the console and server, the printer output may be redirected elsewhere.

Layers

When a frame contains several levels of protocol, the **Detail** view interprets all the levels. The outermost (lowest) level appears first, and so on in order until the innermost (highest) level appears last.

The **Detail** interpretation is often larger than can fit in the available panel on screen. In that case, you see only part of the display at a time, but you may scroll to bring other parts into view.

When you use both **Summary** and **Detail** at the same time, the analyzer prepares a panel for each. The analyzer automatically scrolls the **Detail** view to match the **Summary** view's highlight. When you use **Summary** with **Highest level only**, the server automatically scrolls the **Detail** view so initially it too shows the highest level. When you do not restrict the **Summary** view to **Highest level only**, the analyzer automatically scrolls the **Detail** view to match the level highlighted in the **Summary**.

Controlling the Level Initially Displayed in the Detail View

When you are looking at **Detail** without the **Summary**, you can still control the level to which the analysis server initially scrolls when you switch to a new frame.

To set the level initially displayed in the **Detail** *view*

1. In the Summary display, turn off **Highest level only.**
2. Display frames in the **Summary** view.
3. Move the highlight to the level you want displayed first.
4. Return to **Display Options**. Make the **Summary** view inactive, and **Detail** active.

As you move to a new frame, the analysis server automatically scrolls the **Detail** view to show the level that you last highlighted in the **Summary** view. (When it comes to a frame that doesn't include the level you selected, it shows the next lower level that's actually present.)

Formats for Higher-Level Numeric Addresses

In the **Detail** view, the analyzer shows the numeric form of each source or destination addresses, and also the name you've assigned to address, when it's available in the name table.

The format for numeric display of higher-level addresses is hexadecimal, except where there is an established convention for a different form. For example, a 4-byte IP address is shown as a succession of four decimal numbers separated by dots, with the entire number enclosed in square brackets.

When the analyzer displays an address, it picks a format appropriate to the level and protocol in which the address occurs. Some common examples are shown in Figure 5–13.

<table>
<tr><th></th><th>Token ring or Ethernet</th><th>WAN/
Synchronous</th></tr>
<tr><td><b>DLC</b></td><td>Present in all frames and shown as twelve hexadecimal digits, corresponding to the 6-byte address.
Example: 020701031EF7
Where the Sniffer analyzer recognizes the first three bytes as a manufacturer code, it replaces them by a 6-character abbreviation:
Example: Intrln031EF7</td><td>Every frame is either
<i>from DTE</i> or
<i>from DCE.</i></td></tr>
<tr><td><b>XNS</b></td><td colspan="2">Shown in the same format as 6-byte DLC addresses. (Although an XNS addresses may be the same as a DLC address, the analyzer does not attempt to interpret the manufacturer as it does for DLC addresses.)</td></tr>
<tr><td><b>IP</b></td><td colspan="2">Represented byte by byte, but each byte is represented by its decimal value, a number between 0 and 255. Successive bytes are separated by a dot, and the whole sequence is enclosed in brackets.
Example: [84.12.139.144].</td></tr>
<tr><td><b>DRP
(DECnet)</b></td><td colspan="2">Address represented by two decimal numbers, the area and the node number, separated by a dot. Each number is computed as the binary value after masking certain bits in the address.
Example: 184.27</td></tr>
<tr><td><b>DDP
(AppleTalk)</b></td><td colspan="2">Address represented by two decimal numbers, the network number (representing a 16-bit network address) and the node ID (representing an 8-bit node number), separated by a dot.
Example: 1080.208</td></tr>
</table>

Figure 5–13. Format of address, on selected networks and protocols.

Other address levels and formats may be present, depending upon the optional protocol suites installed in the Sniffer analyzer.

Transmission of 6-Byte DLC Addresses

At the DLC level, every station on earth using 6-byte addressing has a unique address. To be more precise, a DLC address uniquely identifies a station's network interface card. The first three bytes of a DLC address identify the interface card's manufacturer. The IEEE has assigned codes to the various manufacturers. Each manufacturer uses the other three bytes to assign a unique identifier to each of its cards.

Bits on the Wire vs. Bits in Memory

Historically, the various network technologies have had different rules for converting bits "on the wire" to bits in computer memory. Some systems transmit each byte high-order-bit first and some transmit low-order-bit first. Suppose a byte of computer memory contains the value that in hexadecimal is written 87. In memory, that consists of the following bits:

```
1 Ø Ø Ø  Ø 1 1 1
   8        7
```

If you could see the sequence of bits transmitted along a token ring cable, you'd see that each byte is transmitted with the high-order bits first, 1 0 0 0 0 1 1 1. However, if you could make the same observation on Ethernet, you'd see that each byte is transmitted with the low-order bit first. You'd see hex 87 go by as 1 1 1 0 0 0 0 1.

Ordinarily, these different ways of transmitting are completely invisible to any user or program. The interface card translates between bits on the wire and bits in computer memory. You only see bytes in memory, before they've been sent or after they've been received, but never while they're in transit. When an Ethernet card receives the sequence 1 1 1 0 0 0 0 1, it turns that into hex 87, and the receiving station has the same value as the sender. Since (at the DLC level) sender and receiver must be on the same network and must use compatible equipment, a byte in the sender's memory results in the same byte in the receiver's memory. No one at either end need ever know in what sequence the bits within a byte were sent.

Consequence of the IEEE standard

When the IEEE assigned codes to the various manufacturers, it specified the sequence in which bits are transmitted on the wire. It did not specify what the sender or receiver would see in memory. For example, Network General Corporation was assigned a particular sequence of bits. To transmit that sequence on token ring, a Sniffer analysis server has in its memory the bytes 00 00 A6. But to transmit that same sequence on Ethernet, the server's memory must contain the bytes 00 00 65. That's because A6 with the bits of each byte reversed is 65.

Thus, to comply with the IEEE standard, a given manufacturer ID should translate to one number for Ethernet (or any other network using low-order-bit-first), and to a different number for token ring. (At present token ring is the only network using high-order-bit-first.) The tables that the Sniffer analyzer uses to interpret manufacturers' codes make allowance for this. There's one table for token ring and a different table for other networks. Beware, however, that some manufacturers do not follow the letter of IEEE law and use the same value in memory for both network types.

Protocol Interpretation

The Sniffer analyzer gets its ability to interpret protocols from several sources. Interpretation for protocols at the lowest level is included in the software that supports the type of network that the server monitors. Interpreters for higher-level protocols are available in suites. Each suite contains a number of protocols that are related or are likely to occur together. (See the list of protocol interpreter suites in Figure 1–6, page 1–22.)

The interpreters from these different sources are linked into the Sniffer analyzer's executable file when it is built. Each protocol interpreter registers its presence and facilities with the Sniffer analyzer; this permits each interpreter to be represented in the appropriate menus and displays. In the **Detail** view, at the left edge of every line, each interpreter shows an abbreviation indicating the protocol it is decoding.

When you have both **Detail** and **Hex** views open, for each line of the **Detail** interpretation, the corresponding bytes in the **Hex** display are highlighted. In this way you can see together the actual bytes and their interpretation.

While a common registration and display procedure applies to all the interpreters, the various interpreters are nevertheless essentially independent, just as the protocols themselves are essentially independent. Within each protocol, the fields displayed are those that make sense within the context of the protocol.

Bit-Level Interpretation

A protocol may record binary attributes as individual bits packed within a single byte. Where that is done, the interpreters often explode each byte to show its eight bits separately. To illustrate, Figure 5–14 shows the interpretation of hex 1F at a particular position in the Banyan VINES protocol called VIP. Figure 5–15 shows an interpretation of the sequence 6B 80 00 from an IBM SNA header. This sort of decoding interprets not only the meaning of each bit set at a certain position, but also the meaning of each bit *not* set (that is, each 0). Bit-level interpretation is included in almost all protocols; however, in X Windows it is not universally shown, because the number of possible bit-encodings is extremely large.

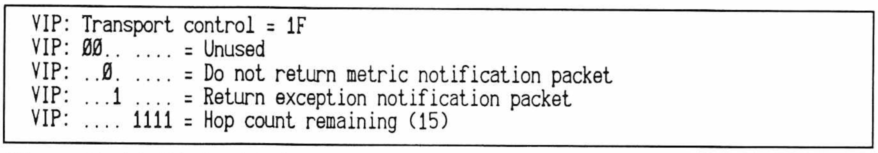

```
VIP: Transport control = 1F
VIP: ØØ.. .... = Unused
VIP: ..Ø. .... = Do not return metric notification packet
VIP: ...1 .... = Return exception notification packet
VIP: .... 1111 = Hop count remaining (15)
```

Figure 5–14. Bit-by-bit interpretation of a byte of VINES internet protocol.

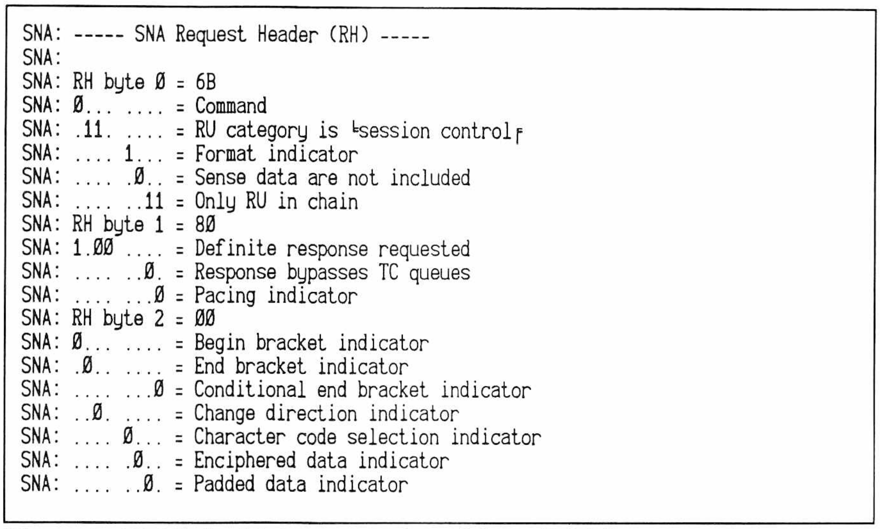

```
SNA: ----- SNA Request Header (RH) -----
SNA:
SNA: RH byte Ø = 6B
SNA: Ø... .... = Command
SNA: .11. .... = RU category is ᴸsession controlᴦ
SNA: .... 1... = Format indicator
SNA: .... .Ø.. = Sense data are not included
SNA: .... ..11 = Only RU in chain
SNA: RH byte 1 = 8Ø
SNA: 1.ØØ .... = Definite response requested
SNA: .... ..Ø. = Response bypasses TC queues
SNA: .... ...Ø = Pacing indicator
SNA: RH byte 2 = ØØ
SNA: Ø... .... = Begin bracket indicator
SNA: .Ø.. .... = End bracket indicator
SNA: .... ...Ø = Conditional end bracket indicator
SNA: ..Ø. .... = Change direction indicator
SNA: .... Ø... = Character code selection indicator
SNA: .... .Ø.. = Enciphered data indicator
SNA: .... ..Ø. = Padded data indicator
```

Figure 5–15. Bit-by-bit interpretation of three SNA bytes on token ring.

Alternate Displays for ASN.1-Encoded Protocols

ISO protocols at the presentation and application levels can be interpreted in two modes. Each is written so that it conforms to a general syntax specified in ASN.1 (Abstract Syntax Notation 1, ISO 8825). The individual protocols then assign meanings to the components identified by the ASN.1 syntax. To accommodate this dual level of interpretation, the ISO protocol interpreter suite gives you the option to interpret frames in these layers in either of two ways:

Syntactic — The Sniffer analyzer labels the ASN.1 components within each frame, or

Semantic — The Sniffer analyzer interprets what the various ASN.1 statements mean, according to the rules of the particular protocol in which the ASN.1 statements occur.

When you select interpretation at the X.400 level, you get the semantic interpretation. To see ASN.1 interpretation, turn off X.400 in the protocol display filter.

Figure 5–16 shows the same fragment of an X.400 frame interpreted first for ASN.1 syntax and then for its content as part of a P1 message envelope.

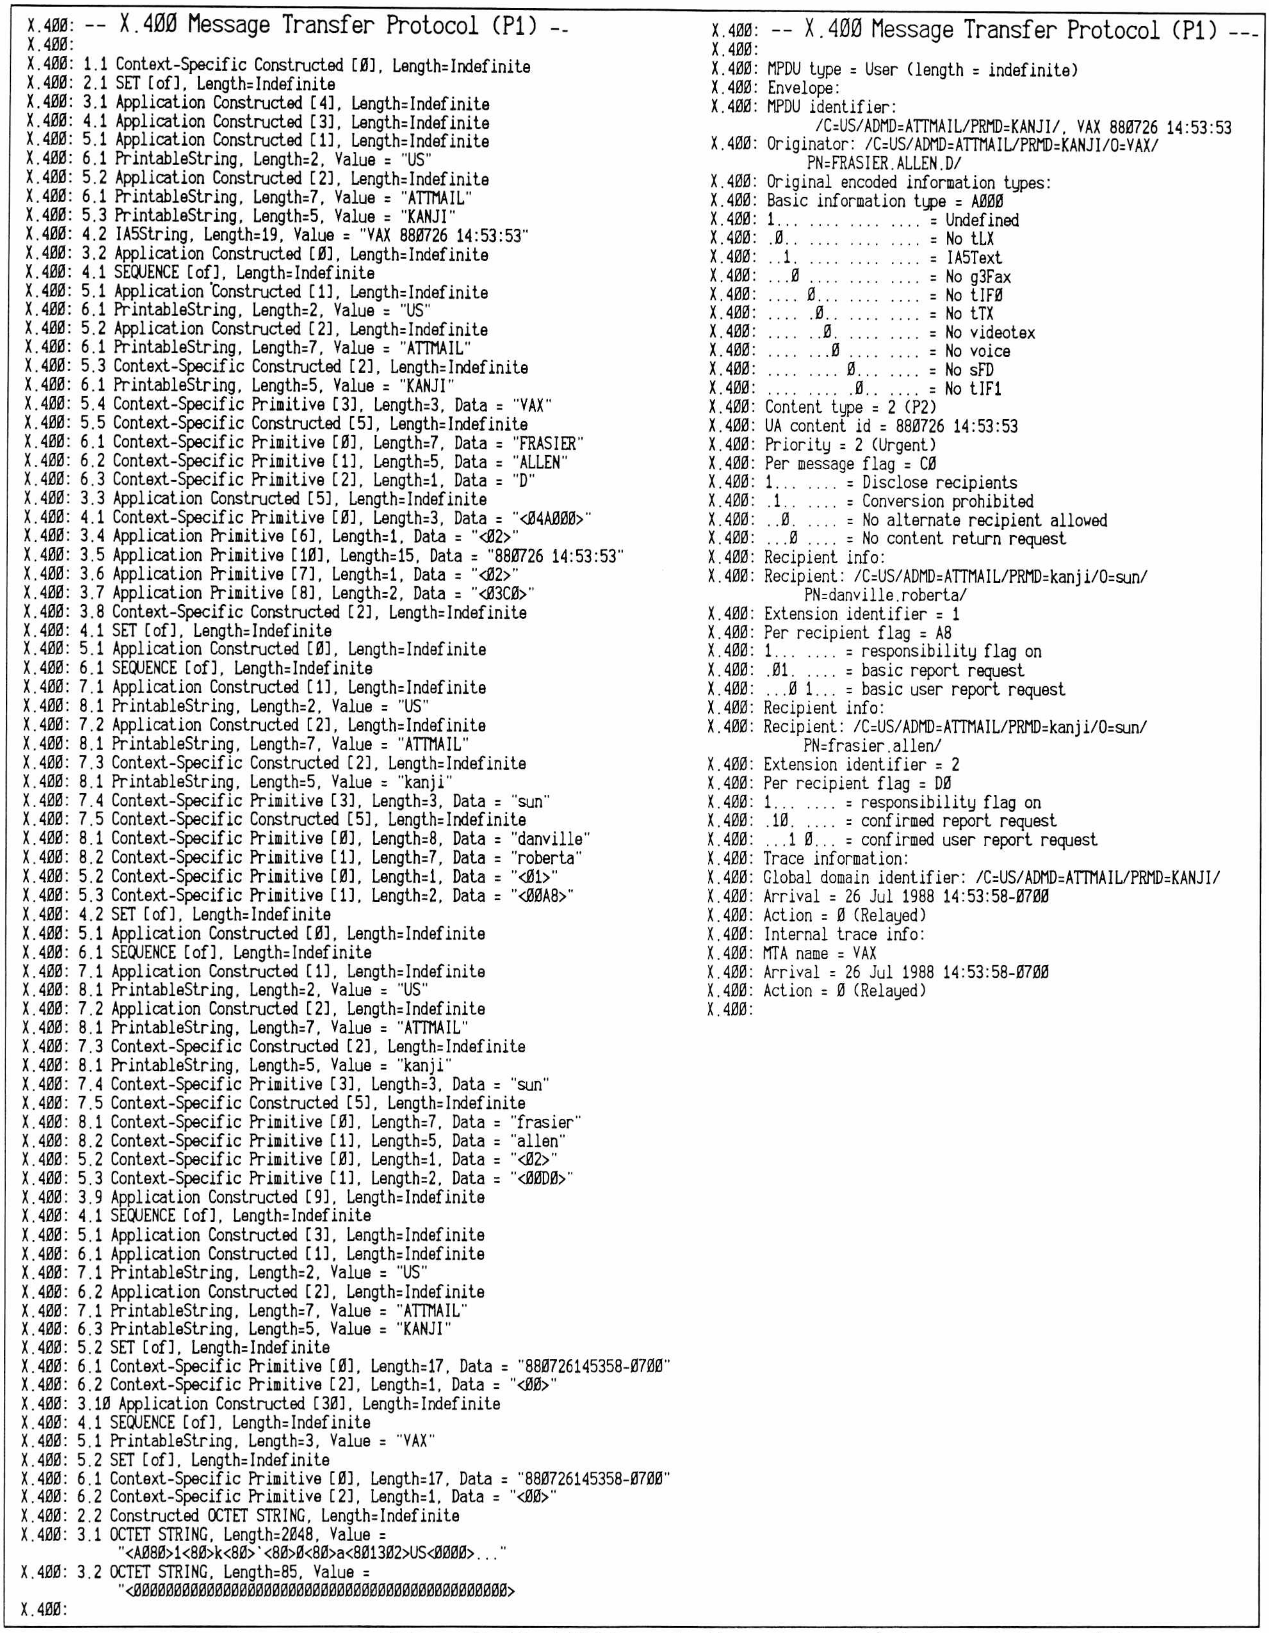

```
X.4ØØ: -- X.4ØØ Message Transfer Protocol (P1) --
X.4ØØ:
X.4ØØ: 1.1 Context-Specific Constructed [Ø], Length=Indefinite
X.4ØØ: 2.1 SET [of], Length=Indefinite
X.4ØØ: 3.1 Application Constructed [4], Length=Indefinite
X.4ØØ: 4.1 Application Constructed [3], Length=Indefinite
X.4ØØ: 5.1 Application Constructed [1], Length=Indefinite
X.4ØØ: 6.1 PrintableString, Length=2, Value = "US"
X.4ØØ: 5.2 Application Constructed [2], Length=Indefinite
X.4ØØ: 6.1 PrintableString, Length=7, Value = "ATTMAIL"
X.4ØØ: 5.3 PrintableString, Length=5, Value = "KANJI"
X.4ØØ: 4.2 IA5String, Length=19, Value = "VAX 88Ø726 14:53:53"
X.4ØØ: 3.2 Application Constructed [Ø], Length=Indefinite
X.4ØØ: 4.1 SEQUENCE [of], Length=Indefinite
X.4ØØ: 5.1 Application Constructed [1], Length=Indefinite
X.4ØØ: 6.1 PrintableString, Length=2, Value = "US"
X.4ØØ: 5.2 Application Constructed [2], Length=Indefinite
X.4ØØ: 6.1 PrintableString, Length=7, Value = "ATTMAIL"
X.4ØØ: 5.3 Context-Specific Constructed [2], Length=Indefinite
X.4ØØ: 6.1 PrintableString, Length=5, Value = "KANJI"
X.4ØØ: 5.4 Context-Specific Primitive [3], Length=3, Data = "VAX"
X.4ØØ: 5.5 Context-Specific Constructed [5], Length=Indefinite
X.4ØØ: 6.1 Context-Specific Primitive [Ø], Length=7, Data = "FRASIER"
X.4ØØ: 6.2 Context-Specific Primitive [1], Length=5, Data = "ALLEN"
X.4ØØ: 6.3 Context-Specific Primitive [2], Length=1, Data = "D"
X.4ØØ: 3.3 Application Constructed [5], Length=Indefinite
X.4ØØ: 4.1 Context-Specific Primitive [Ø], Length=3, Data = "<Ø4AØØØ>"
X.4ØØ: 3.4 Application Primitive [6], Length=1, Data = "<Ø2>"
X.4ØØ: 3.5 Application Primitive [1Ø], Length=15, Data = "88Ø726 14:53:53"
X.4ØØ: 3.6 Application Primitive [7], Length=1, Data = "<Ø2>"
X.4ØØ: 3.7 Application Primitive [8], Length=2, Data = "<Ø3CØ>"
X.4ØØ: 3.8 Context-Specific Constructed [2], Length=Indefinite
X.4ØØ: 4.1 SET [of], Length=Indefinite
X.4ØØ: 5.1 Application Constructed [Ø], Length=Indefinite
X.4ØØ: 6.1 SEQUENCE [of], Length=Indefinite
X.4ØØ: 7.1 Application Constructed [1], Length=Indefinite
X.4ØØ: 8.1 PrintableString, Length=2, Value = "US"
X.4ØØ: 7.2 Application Constructed [2], Length=Indefinite
X.4ØØ: 8.1 PrintableString, Length=7, Value = "ATTMAIL"
X.4ØØ: 7.3 Context-Specific Constructed [2], Length=Indefinite
X.4ØØ: 8.1 PrintableString, Length=5, Value = "kanji"
X.4ØØ: 7.4 Context-Specific Primitive [3], Length=3, Data = "sun"
X.4ØØ: 7.5 Context-Specific Constructed [5], Length=Indefinite
X.4ØØ: 8.1 Context-Specific Primitive [Ø], Length=8, Data = "danville"
X.4ØØ: 8.2 Context-Specific Primitive [1], Length=7, Data = "roberta"
X.4ØØ: 5.2 Context-Specific Primitive [Ø], Length=1, Data = "<Ø1>"
X.4ØØ: 5.3 Context-Specific Primitive [1], Length=2, Data = "<ØØA8>"
X.4ØØ: 4.2 SET [of], Length=Indefinite
X.4ØØ: 5.1 Application Constructed [Ø], Length=Indefinite
X.4ØØ: 6.1 SEQUENCE [of], Length=Indefinite
X.4ØØ: 7.1 Application Constructed [1], Length=Indefinite
X.4ØØ: 8.1 PrintableString, Length=2, Value = "US"
X.4ØØ: 7.2 Application Constructed [2], Length=Indefinite
X.4ØØ: 8.1 PrintableString, Length=7, Value = "ATTMAIL"
X.4ØØ: 7.3 Context-Specific Constructed [2], Length=Indefinite
X.4ØØ: 8.1 PrintableString, Length=5, Value = "kanji"
X.4ØØ: 7.4 Context-Specific Primitive [3], Length=3, Data = "sun"
X.4ØØ: 7.5 Context-Specific Constructed [5], Length=Indefinite
X.4ØØ: 8.1 Context-Specific Primitive [Ø], Length=7, Data = "frasier"
X.4ØØ: 8.2 Context-Specific Primitive [1], Length=5, Data = "allen"
X.4ØØ: 5.2 Context-Specific Primitive [Ø], Length=1, Data = "<Ø2>"
X.4ØØ: 5.3 Context-Specific Primitive [1], Length=2, Data = "<ØØDØ>"
X.4ØØ: 3.9 Application Constructed [9], Length=Indefinite
X.4ØØ: 4.1 SEQUENCE [of], Length=Indefinite
X.4ØØ: 5.1 Application Constructed [3], Length=Indefinite
X.4ØØ: 6.1 Application Constructed [1], Length=Indefinite
X.4ØØ: 7.1 PrintableString, Length=2, Value = "US"
X.4ØØ: 6.2 Application Constructed [2], Length=Indefinite
X.4ØØ: 7.1 PrintableString, Length=7, Value = "ATTMAIL"
X.4ØØ: 6.3 PrintableString, Length=5, Value = "KANJI"
X.4ØØ: 5.2 SET [of], Length=Indefinite
X.4ØØ: 6.1 Context-Specific Primitive [Ø], Length=17, Data = "88Ø726145358-Ø7ØØ"
X.4ØØ: 6.2 Context-Specific Primitive [2], Length=1, Data = "<ØØ>"
X.4ØØ: 3.1Ø Application Constructed [3Ø], Length=Indefinite
X.4ØØ: 4.1 SEQUENCE [of], Length=Indefinite
X.4ØØ: 5.1 PrintableString, Length=3, Value = "VAX"
X.4ØØ: 5.2 SET [of], Length=Indefinite
X.4ØØ: 6.1 Context-Specific Primitive [Ø], Length=17, Data = "88Ø726145358-Ø7ØØ"
X.4ØØ: 6.2 Context-Specific Primitive [2], Length=1, Data = "<ØØ>"
X.4ØØ: 2.2 Constructed OCTET STRING, Length=Indefinite
X.4ØØ: 3.1 OCTET STRING, Length=2Ø48, Value =
           "<AØ8Ø>1<8Ø>k<8Ø>`<8Ø>Ø<8Ø>a<8Ø13Ø2>US<ØØØØ>..."
X.4ØØ: 3.2 OCTET STRING, Length=85, Value =
           "<ØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØØ>
X.4ØØ:
```

```
X.4ØØ: -- X.4ØØ Message Transfer Protocol (P1) ---
X.4ØØ:
X.4ØØ: MPDU type = User (length = indefinite)
X.4ØØ: Envelope:
X.4ØØ: MPDU identifier:
            /C=US/ADMD=ATTMAIL/PRMD=KANJI/, VAX 88Ø726 14:53:53
X.4ØØ: Originator: /C=US/ADMD=ATTMAIL/PRMD=KANJI/O=VAX/
           PN=FRASIER.ALLEN.D/
X.4ØØ: Original encoded information types:
X.4ØØ: Basic information type = AØØØ
X.4ØØ: 1... .... .... .... = Undefined
X.4ØØ: .Ø.. .... .... .... = No tLX
X.4ØØ: ..1. .... .... .... = IA5Text
X.4ØØ: ...Ø .... .... .... = No g3Fax
X.4ØØ: .... Ø... .... .... = No tIFØ
X.4ØØ: .... .Ø.. .... .... = No tTX
X.4ØØ: .... ..Ø. .... .... = No videotex
X.4ØØ: .... ...Ø .... .... = No voice
X.4ØØ: .... .... Ø... .... = No sFD
X.4ØØ: .... .... .Ø.. .... = No tIF1
X.4ØØ: Content type = 2 (P2)
X.4ØØ: UA content id = 88Ø726 14:53:53
X.4ØØ: Priority = 2 (Urgent)
X.4ØØ: Per message flag = CØ
X.4ØØ: 1... .... = Disclose recipients
X.4ØØ: .1.. .... = Conversion prohibited
X.4ØØ: ..Ø. .... = No alternate recipient allowed
X.4ØØ: ...Ø .... = No content return request
X.4ØØ: Recipient info:
X.4ØØ: Recipient: /C=US/ADMD=ATTMAIL/PRMD=kanji/O=sun/
           PN=danville.roberta/
X.4ØØ: Extension identifier = 1
X.4ØØ: Per recipient flag = A8
X.4ØØ: 1... .... = responsibility flag on
X.4ØØ: .Ø1. .... = basic report request
X.4ØØ: ...Ø 1... = basic user report request
X.4ØØ: Recipient info:
X.4ØØ: Recipient: /C=US/ADMD=ATTMAIL/PRMD=kanji/O=sun/
           PN=frasier.allen/
X.4ØØ: Extension identifier = 2
X.4ØØ: Per recipient flag = DØ
X.4ØØ: 1... .... = responsibility flag on
X.4ØØ: .1Ø. .... = confirmed report request
X.4ØØ: ...1 Ø... = confirmed user report request
X.4ØØ: Trace information:
X.4ØØ: Global domain identifier: /C=US/ADMD=ATTMAIL/PRMD=KANJI/
X.4ØØ: Arrival = 26 Jul 1988 14:53:58-Ø7ØØ
X.4ØØ: Action = Ø (Relayed)
X.4ØØ: Internal trace info:
X.4ØØ: MTA name = VAX
X.4ØØ: Arrival = 26 Jul 1988 14:53:58-Ø7ØØ
X.4ØØ: Action = Ø (Relayed)
X.4ØØ:
```

Figure 5–16. ASN.1 syntactic and semantic interpretations of the same X.400 layer of an ISO frame.

Token Ring "Address Recognized" and "Frame Copied" Bits

As each frame circulates on the token ring, it has two extra "trailer" bytes at the end. The trailer is used to check the frame's validity. It isn't usually considered part of the frame, and doesn't appear in the **Hex** view. However, the **Detail** view reports the status of two indicators in the trailer: the address recognized bit and the frame copied bit. They're visible in Figure 5–17 as part of the **Detail** view's DLC report.

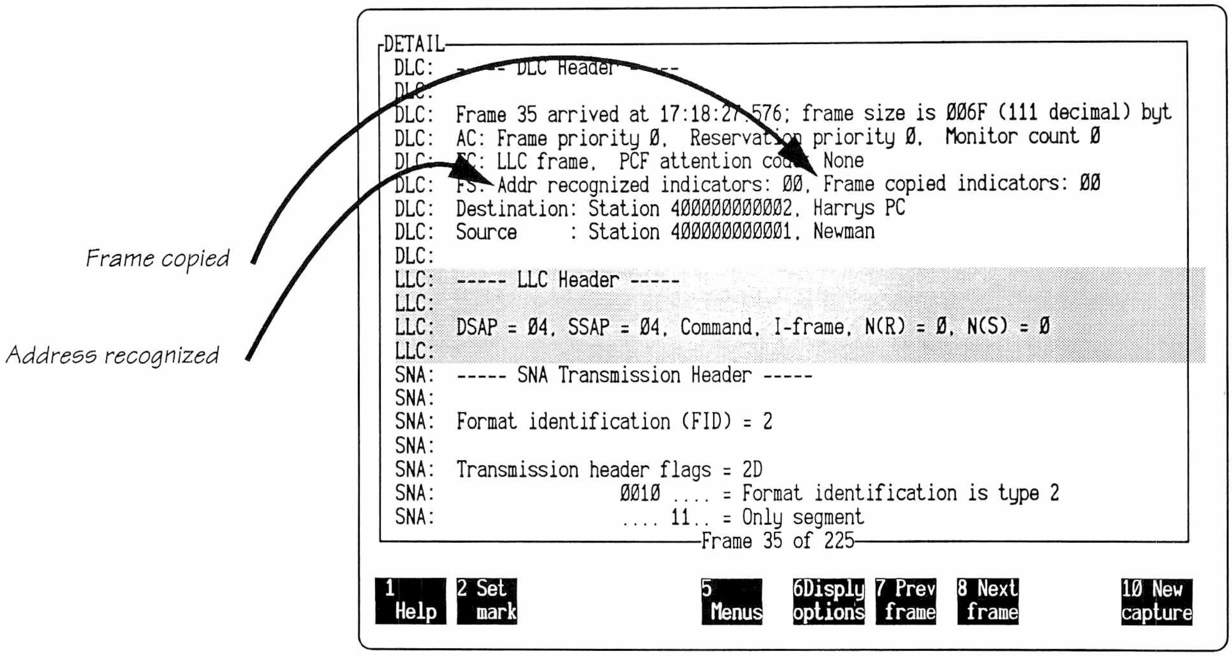

```
DETAIL
 DLC:  ----- DLC Header -----
 DLC:
 DLC:  Frame 35 arrived at 17:18:27.576; frame size is ØØ6F (111 decimal) byt
 DLC:  AC: Frame priority Ø, Reservation priority Ø, Monitor count Ø
 DLC:  FC: LLC frame, PCF attention code: None
 DLC:  FS: Addr recognized indicators: ØØ, Frame copied indicators: ØØ
 DLC:  Destination: Station 4ØØØØØØØØØØ2, Harrys PC
 DLC:  Source     : Station 4ØØØØØØØØØØ1, Newman
 DLC:
 LLC:  ----- LLC Header -----
 LLC:
 LLC:  DSAP = Ø4, SSAP = Ø4, Command, I-frame, N(R) = Ø, N(S) = Ø
 LLC:
 SNA:  ----- SNA Transmission Header -----
 SNA:
 SNA:  Format identification (FID) = 2
 SNA:
 SNA:  Transmission header flags = 2D
 SNA:                    ØØ1Ø .... = Format identification is type 2
 SNA:                    .... 11.. = Only segment
                             Frame 35 of 225
```

Figure 5–17. Detail view showing token ring "address recognized" and "frame copied" bits.

When a station recognizes itself in the frame's destination, it sets the *address recognized* bit. If that bit is on when the frame reaches you, at least one station upstream from you has recognized itself as a recipient. (A frame sent to a functional address may have any number of recipients.)

When a station retains a copy of the frame, it sets the *frame copied* bit. Normally, a recipient both recognizes and records the frame, and sets both bits.

The values you see for *addressed recognized* and *frame copied* depend on where you are located. If a frame reaches you with *address recognized* already set, the frame must have reached the recipient before it reached you. When you're looking for problems at a particular portion of the ring, you may wish to capture from different positions. The *address recognized* bit provides evidence of your position in the ring sequence. To verify that a station has accepted frames addressed to it, you have to be upstream from the sender and downstream from the recipient. (That is, you must not be between the sender and recipient in ring order.)

Use of Color

Except when the SniffMaster console lacks a color display, each layer of protocol has its own characteristic shade.

Where possible, each layer is identified with one of the seven layers of the OSI model. (However, many protocols predate the OSI model, and don't fit into it very clearly.) The Sniffer analyzer assigns colors to the layers as shown in Figure 5–18.

Protocol	Layer	Color
Physical level protocols	1	Magenta
Fragmentation protocols	1+	Red
Link level protocols	2	Brown
Network level protocols	3	Green
Transport level protocols	4	Yellow
Session level protocols	5	Light green
Presentation level protocols	6	Light cyan
Application level protocols I	7–	Light red
Application level protocols II	7	Light magenta
Name protocols and network management layers		Cyan (blue-green)
Various protocol glue layers		Black
Other		Light blue
Uninterpreted data		Gray

Figure 5–18. Foreground colors to layers of the OSI model.

Normal color displays use a blue background. Highlighted areas have a light-blue background. However, highlighted portions of a hexadecimal display have white backgrounds.

The Hexadecimal View

The **Hex** view displays each byte as two hex characters, 00 to FF, with a blank between successive bytes. The bytes are arranged 16 to a row in a full-width table (8 to a row in the half-width table for two viewports). At the left, the offset from the beginning of the frame is displayed in hexadecimal (Figure 5–19).

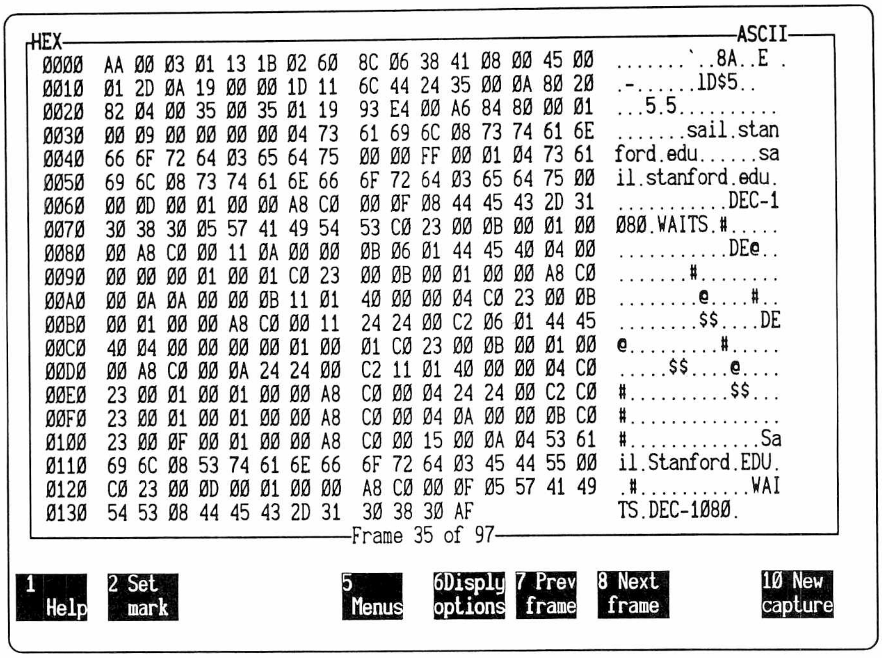

```
HEX                                                                ASCII
 ØØØØ  AA ØØ Ø3 Ø1 13 1B Ø2 6Ø  8C Ø6 38 41 Ø8 ØØ 45 ØØ  .......`..8A..E .
 ØØ1Ø  Ø1 2D ØA 19 ØØ ØØ 1D 11  6C 44 24 35 ØØ ØA 8Ø 2Ø  .-......lD$5..
 ØØ2Ø  82 Ø4 ØØ 35 ØØ 35 Ø1 19  93 E4 ØØ A6 84 8Ø ØØ Ø1  ...5.5..........
 ØØ3Ø  ØØ Ø9 ØØ ØØ ØØ ØØ Ø4 73  61 69 6C Ø8 73 74 61 6E  .......sail.stan
 ØØ4Ø  66 6F 72 64 Ø3 65 64 75  ØØ ØØ FF ØØ Ø1 Ø4 73 61  ford.edu......sa
 ØØ5Ø  69 6C Ø8 73 74 61 6E 66  6F 72 64 Ø3 65 64 75 ØØ  il.stanford.edu.
 ØØ6Ø  ØØ ØD ØØ Ø1 ØØ ØØ A8 CØ  ØØ ØF Ø8 44 45 43 2D 31  ...........DEC-1
 ØØ7Ø  3Ø 38 3Ø Ø5 57 41 49 54  53 CØ 23 ØØ ØB ØØ Ø1 ØØ  Ø8Ø.WAITS.#.....
 ØØ8Ø  ØØ A8 CØ ØØ 11 ØA ØØ ØØ  ØB Ø6 Ø1 44 45 4Ø Ø4 ØØ  ...........DE@..
 ØØ9Ø  ØØ ØØ ØØ Ø1 ØØ Ø1 CØ 23  ØØ ØB ØØ Ø1 ØØ ØØ A8 CØ  .......#........
 ØØAØ  ØØ ØA ØA ØØ ØØ ØB 11 Ø1  4Ø ØØ ØØ Ø4 CØ 23 ØØ ØB  ........@....#..
 ØØBØ  ØØ Ø1 ØØ ØØ A8 CØ ØØ 11  24 24 ØØ C2 Ø6 Ø1 44 45  ........$$....DE
 ØØCØ  4Ø Ø4 ØØ ØØ ØØ ØØ Ø1 ØØ  Ø1 CØ 23 ØØ ØB ØØ Ø1 ØØ  @.........#.....
 ØØDØ  ØØ A8 CØ ØØ ØA 24 24 ØØ  C2 11 Ø1 4Ø ØØ ØØ Ø4 CØ  .....$$....@....
 ØØEØ  23 ØØ Ø1 ØØ Ø1 ØØ ØØ A8  CØ ØØ Ø4 24 24 ØØ C2 CØ  #..........$$...
 ØØFØ  23 ØØ Ø1 ØØ Ø1 ØØ ØØ A8  CØ ØØ Ø4 ØA ØØ ØØ ØB CØ  #...............
 Ø1ØØ  23 ØØ ØF ØØ Ø1 ØØ ØØ A8  CØ ØØ 15 ØØ ØA Ø4 53 61  #.............Sa
 Ø11Ø  69 6C Ø8 53 74 61 6E 66  6F 72 64 Ø3 45 44 55 ØØ  il.Stanford.EDU.
 Ø12Ø  CØ 23 ØØ ØD ØØ Ø1 ØØ ØØ  A8 CØ ØØ ØF Ø5 57 41 49  .#...........WAI
 Ø13Ø  54 53 Ø8 44 45 43 2D 31  3Ø 38 3Ø AF              TS.DEC-1Ø8Ø.
                                  Frame 35 of 97
```

Figure 5–19. Hexadecimal view of a TCP/IP frame on Ethernet.

Text Transliteration

To the right of the hexadecimal codes, the **Hex** view shows the corresponding ASCII or EBCDIC characters. A standard character is shown by its text equivalent; anything else is represented by a dot.

The transliteration of characters follows either ASCII or EBCDIC conventions. For Ethernet, the **Hex** menu includes a radio-control with two choices:

▶ ASCII characters
EBCDIC characters

The default is ASCII. The choice applies to all levels of all frames.

To select ASCII or EBCDIC transliteration

1. In the main menu, select **Display**, then **Hex**, then move to the panel to the right of **Hex**.
2. Move the highlight to **ASCII** or **EBCDIC** and press Spacebar. That moves the ▶ to the highlighted line.

Dynamic Choice of Transliteration

On token ring or WAN/synchronous link, the menu includes a third choice: **Dynamic**. It's the default.

ASCII characters
EBCDIC characters
▶ Dynamic mode

WAN/Synchronous Dynamic Transliteration

When you select **Dynamic mode**, the transliteration depends on the packet type set in the **Options** menu. When **packet type** is SDLC/ SNA, transliteration is EBCDIC. Otherwise, it's ASCII.

Token Ring Dynamic Transliteration

When you select **Dynamic mode**, the Sniffer analyzer decides separately for each frame whether the frame should be transliterated as ASCII or EBCDIC. It gives you EBCDIC display of all MAC frames, and of any LLC frame whose SAP indicates that it contains SNA. It gives ASCII display of all other frames.

Windows, Views, and Scrolling

When you display two or three views at once, panels in the active viewport scroll together. For example, when you're displaying both **Summary** and **Detail**, you might scroll the **Summary** view so that it shows the next frame. That causes the **Detail** view to show the next frame also. The views remain in step.

Similarly, when you shift the highlight in the **Summary** view to mark a different frame, the viewport's **Detail** view and its **Hex** view both scroll automatically, so that they show the frame highlighted in the **Summary** view.

When you're showing multiple levels within the **Summary** view, the report on a single frame may occupy several lines (one for each level). You can move the highlight from one level to the next, yet remain within the same frame. For example, you might move the highlight from the DLC level to the level above. When you do that, the **Detail** view also scrolls. The level at the top of the **Detail** view matches the level you've highlighted in the **Summary** view.

Highlighting Detail in the Hex View

When you display both **Detail** and **Hex** views, as you move the highlight in the **Detail** view, the Sniffer analyzer highlights the corresponding bytes in the **Hex** view (Figure 5–20). That makes it easy to match a sequence of bytes with its interpretation.

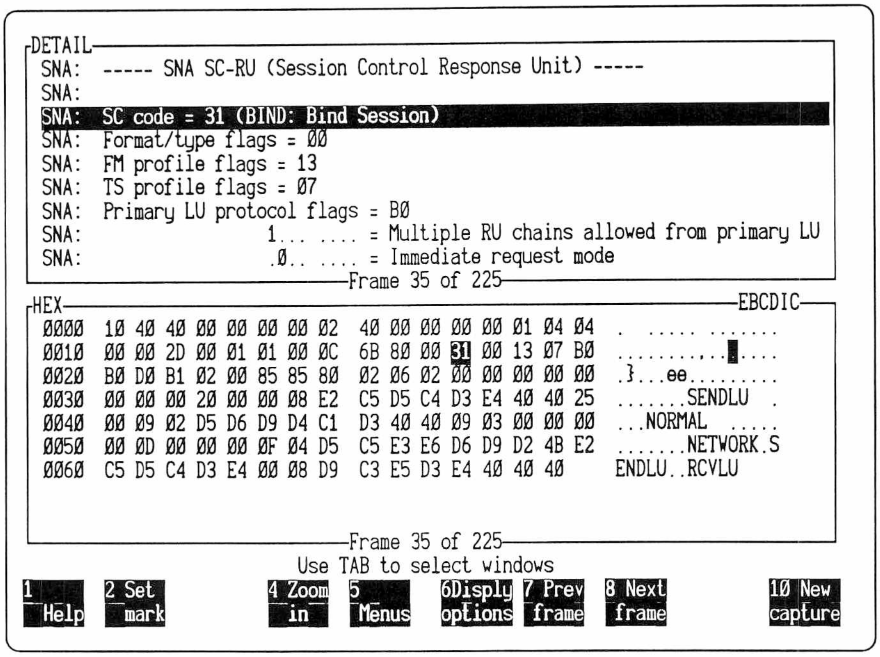

```
DETAIL
 SNA:  ----- SNA SC-RU (Session Control Response Unit) -----
 SNA:
 SNA:  SC code = 31 (BIND: Bind Session)
 SNA:  Format/type flags = ØØ
 SNA:  FM profile flags = 13
 SNA:  TS profile flags = Ø7
 SNA:  Primary LU protocol flags = BØ
 SNA:                 1... .... = Multiple RU chains allowed from primary LU
 SNA:                 .Ø.. .... = Immediate request mode
                         Frame 35 of 225
HEX                                                                EBCDIC
 ØØØØ  1Ø 4Ø 4Ø ØØ ØØ ØØ ØØ Ø2  4Ø ØØ ØØ ØØ ØØ Ø1 Ø4 Ø4  .  ..... .......
 ØØ1Ø  ØØ ØØ 2D ØØ Ø1 Ø1 ØØ ØC  6B 8Ø ØØ 31 ØØ 13 Ø7 BØ  ..........█....
 ØØ2Ø  BØ DØ B1 Ø2 ØØ 85 85 8Ø  Ø2 Ø6 Ø2 ØØ ØØ ØØ ØØ ØØ  .}...ee.........
 ØØ3Ø  ØØ ØØ ØØ 2Ø ØØ ØØ Ø8 E2  C5 D5 C4 D3 E4 4Ø 4Ø 25  .......SENDLU  .
 ØØ4Ø  ØØ Ø9 Ø2 D5 D6 D9 D4 C1  D3 4Ø 4Ø Ø9 Ø3 ØØ ØØ ØØ  ...NORMAL  .....
 ØØ5Ø  ØØ ØD ØØ ØØ ØØ ØF Ø4 D5  C5 E3 E6 D6 D9 D2 4B E2  .......NETWORK.S
 ØØ6Ø  C5 D5 C4 D3 E4 ØØ Ø8 D9  C3 E5 D3 E4 4Ø 4Ø 4Ø     ENDLU..RCVLU

                         Frame 35 of 225
                     Use TAB to select windows
```

Figure 5–20. Synchronized highlighting in the **Detail** *and* **Hex** *views.*

The **Hex** view shows the offset to the start of each line. You can readily calculate each field's address. The hexadecimal offset is required when you describe a pattern to be matched. However, you can use the Cursor Up key to "cut and paste" from a displayed frame, without having to type the hex offset yourself.

Detail and Hex Display of Spanned Frames

In some protocols, a message at one of the higher levels may span several DLC frames. (See Figure 2–24, page 2–38.) During display, the Sniffer analyzer's **Detail** view reassembles the entire high-level message as if the whole message were in the first frame. You can read the message continuously without having to jump among the separate frames that contain its parts. When you highlight a part of the **Detail** display, the **Hex** view continues to highlight the corresponding hex characters.

Figure 5–21 illustrates the treatment of a spanned high-level message. Here are some points to observe:

- In the summary view, the frame in which the high-level message starts has a note telling you how many frames the message spans. In Figure 5–21, the note says ISO_TP Data EOT (5 frames). In this example, the ISO TP data was split among frames 15, 16, 17, 18, 19 and 20. There's no presumption that spanned frames are consecutive; unrelated frames might have arrived between them and so appear interspersed in the display.

- In the **Detail** view, the entire high-level message is displayed as though all of it were in the frame in which it starts. When you print the display, the entire high-level message appears without a break, taking as many lines as necessary. On screen, you can see the rest of the message by scrolling within the **Detail** view.

- Scrolling within the **Detail** view (as usual) causes the **Hex** view to scroll automatically. The field you highlight has a corresponding highlight in the **Hex** view.

- The frame number in the **Hex** view doesn't necessarily match the frame number in the **Detail** view. In Figure 5–21, the interpretation of ISO Session Layer is part of the **Detail** view of frame 15, because that's where the ISO TP data starts. However, the corresponding **Hex** view shows frame 17. That's because (although the message started in frame 15) its ISO Session Layer is in frame 17, starting at offset 36.

- When the field highlighted in the **Detail** view extends over additional frames, the **Hex** panel scrolls to the start of the field, but adds a + sign to show that there's more present in a different frame. A plus is visible in the last row of the **Hex** view in Figure 5–21.

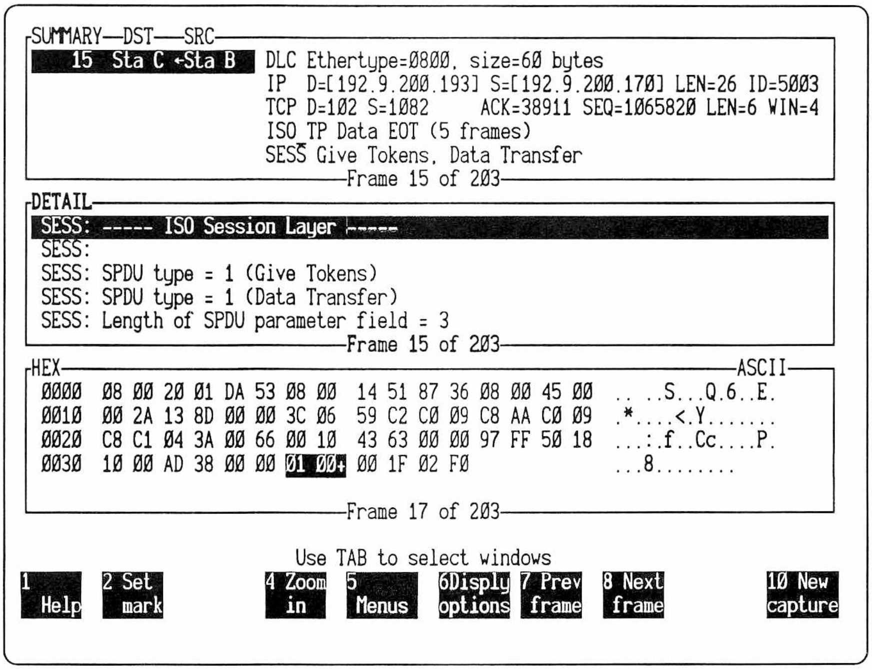

```
┌SUMMARY—DST——SRC——————————————————————————————————————————————————————————
  15  Sta C ←Sta B  DLC Ethertype=Ø8ØØ, size=6Ø bytes
                    IP  D=[192.9.2ØØ.193] S=[192.9.2ØØ.17Ø] LEN=26 ID=5ØØ3
                    TCP D=1Ø2 S=1Ø82     ACK=38911 SEQ=1Ø6582Ø LEN=6 WIN=4
                    ISO TP Data EOT (5 frames)
                    SESS Give Tokens, Data Transfer
                    ————————————Frame 15 of 2Ø3———————————————————————————
┌DETAIL————————————————————————————————————————————————————————————————————
 SESS: ----- ISO Session Layer -----
 SESS:
 SESS: SPDU type = 1 (Give Tokens)
 SESS: SPDU type = 1 (Data Transfer)
 SESS: Length of SPDU parameter field = 3
                    ————————————Frame 15 of 2Ø3———————————————————————————
┌HEX————————————————————————————————————————————————————————————ASCII———————
 ØØØØ  Ø8 ØØ 2Ø Ø1 DA 53 Ø8 ØØ  14 51 87 36 Ø8 ØØ 45 ØØ  .. ..S...Q.6..E.
 ØØ1Ø  ØØ 2A 13 8D ØØ ØØ 3C Ø6  59 C2 CØ Ø9 C8 AA CØ Ø9  .*....<.Y.......
 ØØ2Ø  C8 C1 Ø4 3A ØØ 66 ØØ 1Ø  43 63 ØØ ØØ 97 FF 5Ø 18  ...:.f..Cc....P.
 ØØ3Ø  1Ø ØØ AD 38 ØØ ØØ Ø1 ØØ+ ØØ 1F Ø2 FØ              ...8........

                    ————————————Frame 17 of 2Ø3———————————————————————————

                           Use TAB to select windows
```

Figure 5–21. **Detail** *and* **Hex** *display of a spanned ISO frame.*

Function Keys Available During Display

During display, the following function keys are active:

F1 **Help**. Provides assistance in using the Sniffer analyzer.

F2 **Set mark.** Marks the reference frame with an M. The values reported by relative time and cumulative bytes count from the marked frame (see page 5–22).

F3 **Display**. When display has been interrupted (for example, by pressing F6 to change the display options), F3 is enabled. Pressing it resumes display.

F4 **Zoom in/out.** Temporarily expands the active panel to fill the entire window. When several panels are open at once, the space for an individual panel may be too small to show much. Pressing F4 zooms so that the active panel has the entire window. The other panels are concealed. Pressing F4 a second time restores the tiled arrangement of the open panels.

Zooming enlarges all the panels, but places them one above another so you only see one at a time. You can still use the Tab key to move to the next open panel, which then gets the whole window.

F5 **Menus.** Returns you to the main menu.

F6 **Display options.** Provides you with the basic set of display options found in the Sniffer analyzer's main menu. It also offers several auxiliary options, permitting you to:

— Go directly to a specific frame number.

— Search for a text string.

— Search for a pattern in a frame.

— Jump directly to the marked frame.

— Jump directly to the trigger frame.

See the following section, Searching and Jumping, for further information on these choices.

F7 **Previous frame**. Takes you to the previous frame. Only frames that pass the current display filter are shown, so pressing this key takes you to the adjacent visible frame, perhaps skipping over one or more that the filter has rejected.

In a summary view that shows several frames at once, F7 (or the Cursor Up) moves to the preceding line (either the preceding frame or the preceding level of the current frame). It does so by moving the zone that is highlighted and scrolling if the highlight is already at the top or bottom of the panel.

In the **Detail** view and the **Hex** view, F7 replaces the current display with the display for the preceding frame.

If you have multiple views within a viewport —that is, summary, **Detail** and hexadecimal— when scrolling takes you to another frame, your hexadecimal or **Detail** views (if open) move to that frame, too.

F8 **Next Frame.** Pressing F8 (or the Cursor Down key) takes you to the next line or the next frame in the same way that F7 takes you to the preceding line or frame.

F10 **New Capture.** Starts the capture process again.

Cursor movement keys

↑ Scrolls to the preceding line of the display.

↓ Scrolls to the following line of the display

← Scrolls the display so 8 columns to the left are brought into view.

→ Scrolls the display so 8 columns to the right are brought into view.

PgUp Scrolls to the preceding page of the display.

PgDown Scrolls to the following page of the display.

Home Jumps to the top of the display.

End Jumps to the bottom of the display.

Tab Moves to next panel.

Shift Tab Moves to previous panel.

Searching and Jumping

When you press F6 during display, you get a menu of display options superimposed on your current display. It includes several ways of searching and moving rapidly to particular frames (Figure 5–22).

Jump to mark Jumps to the frame marked M. (If you haven't yet marked a frame, the *first* frame is marked.)

Jump to trigger Jumps to the frame marked T.

Go to Frame Moves directly to a frame whose number you know.

To go to a frame identified by number

1. During display, press F6 to overlay the **Display Options** menu.
2. Move the highlight to **Go to frame nn** and press Enter.

Result: The analyzer opens an additional dialog box in which you can write the number of the frame you want. It won't let you ask for a frame whose number is larger than the number of frames in the buffer.

3. Type the frame number and Press Enter.

   Result: The analyzer removes the overlaid dialog box and resumes display at the frame you requested.

Search for Text

One of the options available during display is to search through the capture buffer for a frame that contains a particular text. The text doesn't have to be physically present in the frame. You can search for text that exists only in the analyzer's interpretation of the frame.

To search the capture buffer for a frame containing particular text

1. During display, press F6 to overlay the **Display Options** menu.
2. Move the highlight to **Search for text** (Figure 5–22). Don't press Enter yet (you haven't yet told it what to search for). Move to the panel to the right.

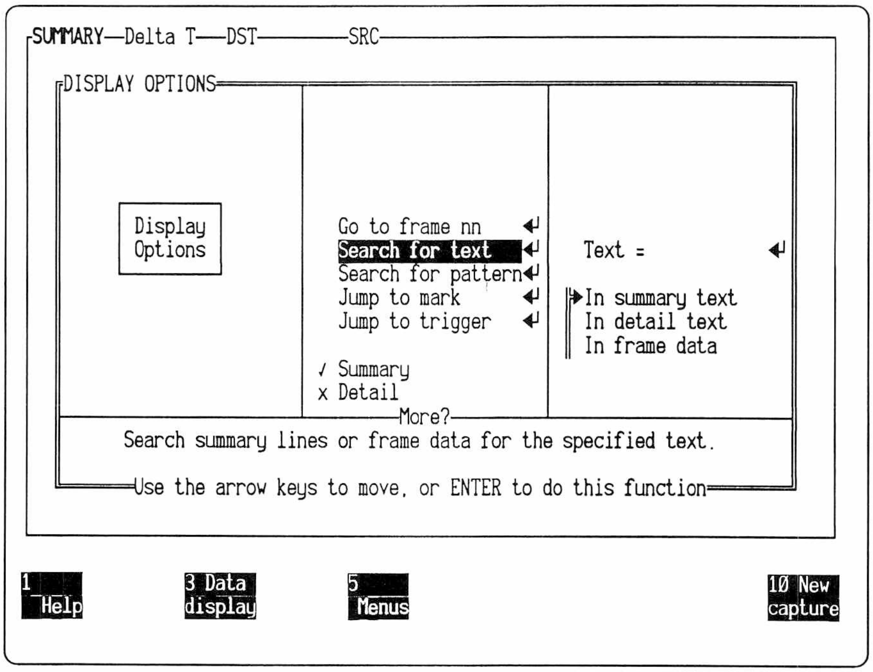

Search summary lines or frame data for the specified text.

Use the arrow keys to move, or ENTER to do this function

Figure 5–22. Menu for requesting a text search.

3. Move the highlight to **Text =**, and press Enter.

Result: The analyzer opens a dialog box headed **Enter text.** Type up to 31 characters. The search is case sensitive. Press Enter to record the text you want to look for (Figure 5–23).

4. Then tell the analyzer where to look—that is, whether to look in the **Summary** text, the **Detail** text, or the text of the frame data. Move the highlight to the line you want and press Spacebar to move the arrow of the radio control.

   — (When you choose **Frame data** as the place to look, you have the further choice of searching the text's ASCII or EBCDIC representation.)

5. Move the highlight back to **search for text**, and press Enter.

   Result: The Sniffer analyzer moves forward through the frames, searching for the text you entered. When it completes its search, it removes the overlaid dialog box and resumes display at the frame containing the target text.

When you ask for a search in the **Summary** or **Detail** view, the search is not for characters in the frame itself, but for characters that appear in the analysis server's interpretation. The analysis server generates the **Summary, Detail,** or **Hex** display for each frame and searches it for characters that match your request.

Figure 5–23 shows a request to search some LocalTalk frames for the phrase "Distance = 6" in the **Detail** view.

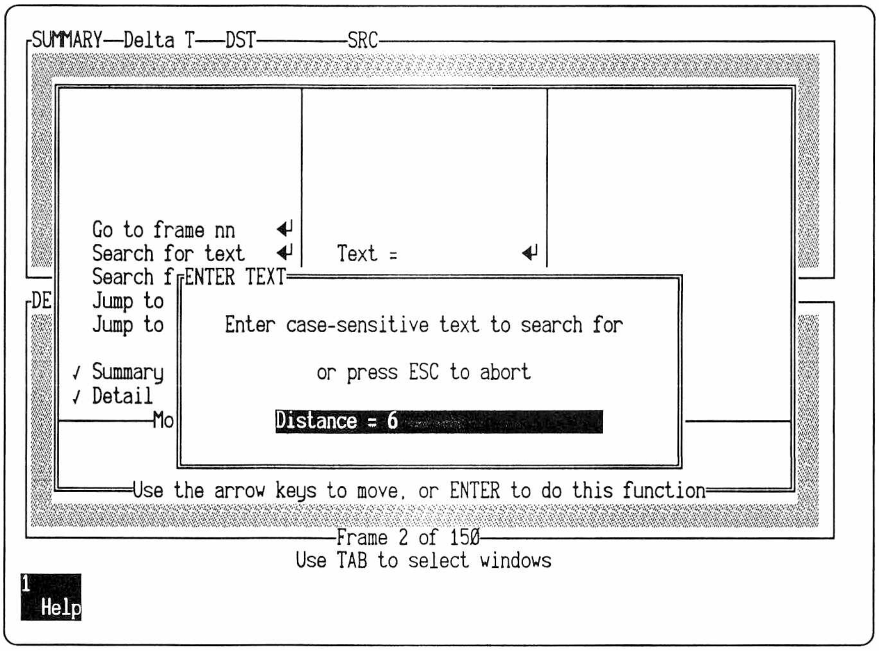

```
SUMMARY—Delta T—DST————SRC——————————————

  Go to frame nn     ↵
  Search for text    ↵   Text =           ↵
  Search f ENTER TEXT
DE Jump to
  Jump to     Enter case-sensitive text to search for

 √ Summary            or press ESC to abort
 √ Detail
 ————Mo          Distance = 6

 ———Use the arrow keys to move, or ENTER to do this function———
                    —Frame 2 of 150—
                    Use TAB to select windows
1
 Help
```

Figure 5–23. Dialog box to enter text to search for.

The search starts with the frame following the one you are looking at and stops at the first match. If the search reaches the last frame in the capture buffer without finding a match, searching continues from the first frame. If all frames have been searched without finding a match, the analysis server reports that and stops searching.

To repeat a search for text

1. Press F6 again (to return to the display options menu).
2. Press S (for search for text).

   Result: The text you last searched for is still in the field marked **Text =.**
3. If you're searching again for the same text, just press Enter. Otherwise, replace the search text and then press Enter.

Search for Pattern

You can search for a frame containing a particular pattern. Figure 5–24 shows the menu for specifying a search pattern.

To search the capture buffer for a pattern

1. Press F6 again (to return to the display options menu).
2. Move the highlight to **Search for Pattern** (Figure 5–24).
3. Move to the right to specify the pattern.

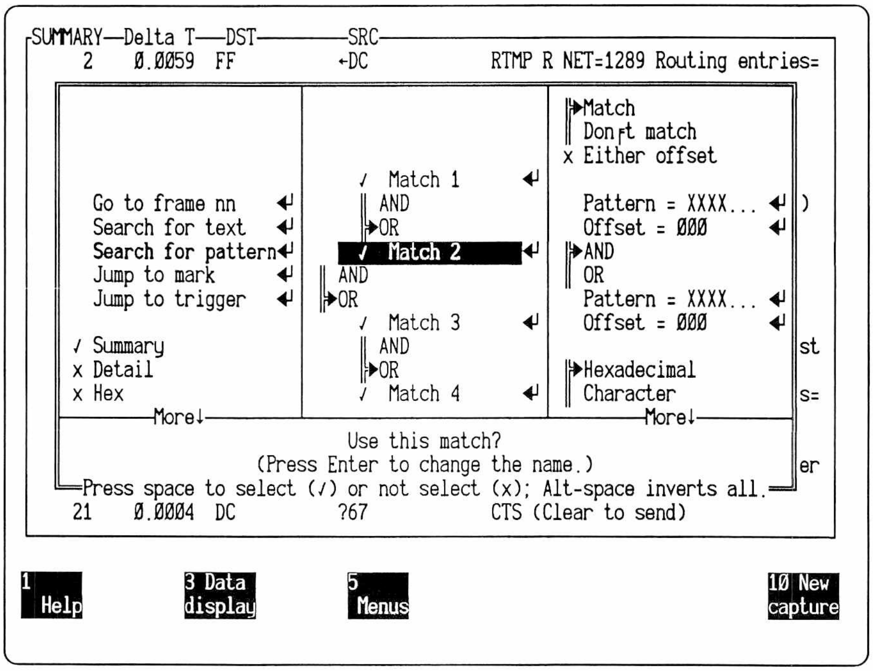

```
SUMMARY—Delta T—DST————SRC——————————————————————————
  2   Ø.ØØ59  FF         ←DC        RTMP R NET=1289 Routing entries=

                                              ▶Match
                                               Don't match
                                             x Either offset
                             ✓ Match 1    ↵
   Go to frame nn     ↵        AND               Pattern = XXXX...  ↵  )
   Search for text    ↵       ▶OR                Offset = ØØØ       ↵
   Search for pattern ↵      ✓ Match 2    ↵   ▶AND
   Jump to mark       ↵    AND                 OR
   Jump to trigger    ↵   ▶OR                    Pattern = XXXX...  ↵
                             ✓ Match 3    ↵     Offset = ØØØ       ↵
 ✓ Summary                     AND                                    st
 x Detail                     ▶OR              ▶Hexadecimal
 x Hex                       ✓ Match 4    ↵     Character             s=
           More↓                                        More↓
                          Use this match?
                   (Press Enter to change the name.)                  er
 Press space to select (✓) or not select (x); Alt-space inverts all.
 21   Ø.ØØØ4  DC          ?67          CTS (Clear to send)

1           3 Data       5                                10 New
  Help      display       Menus                           capture
```

Figure 5–24. Specifying a pattern to search for.

The procedure is exactly the same as the one for specifying a pattern in the capture filter (see page 2–19ff).

4. After you've specified the pattern, return to **Search for Pattern** and press Enter.

   Result: The Sniffer analyzer moves forward through the frames, searching for the pattern you entered. When it completes its search, it removes the overlaid dialog box and resumes display at the frame containing the target text.

As with text search, the search starts with the frame following the one you are looking at and stops at the first match. If the search reaches the last frame in the capture buffer without finding a match, searching continues from the first frame. If all frames have been searched without finding a match, the analysis server reports that and stops searching.

When the Sniffer analyzer finds the pattern you requested, it displays the message
`Found at frame` *nn*.
If the **Summary** view is open, it highlights the frame.

To repeat a search for pattern

1. Press F6 again (to return to the display options menu)
2. Move the highlight to **Search for Pattern**.

   Result: The pattern you last searched for is still there.

3. If you're searching again for the same pattern, just press Enter. Otherwise, revise the pattern, then return to **Search for pattern** and press Enter.

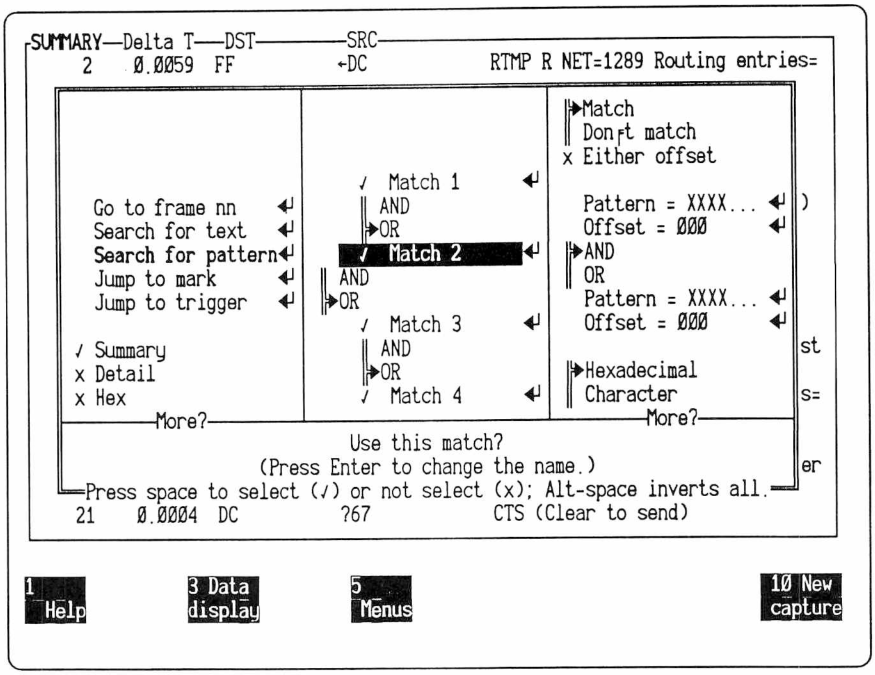

```
SUMMARY—Delta T—DST————SRC————————————————————————————————————
  2    Ø.ØØ59  FF          ←DC            RTMP R NET=1289 Routing entries=
                                                  ►Match
                                                   Don't match
                                                  x Either offset
                                  ✓ Match 1    ↵
   Go to frame nn    ↵             AND              Pattern = XXXX... ↵ )
   Search for text   ↵            ►OR               Offset = ØØØ      ↵
   Search for pattern↵         ✓ Match 2       ↵   ►AND
   Jump to mark      ↵         AND                  OR
   Jump to trigger   ↵        ►OR                   Pattern = XXXX... ↵
                                  ✓ Match 3    ↵    Offset = ØØØ      ↵
 ✓ Summary                         AND                                    st
 x Detail                         ►OR              ►Hexadecimal
 x Hex                            ✓ Match 4    ↵    Character             S=
          More?                                           More?
                          Use this match?
                 (Press Enter to change the name.)                        er
 Press space to select (✓) or not select (x); Alt-space inverts all.
  21    Ø.ØØØ4  DC          ?67           CTS (Clear to send)

1 Help      3 Data display      5 Menus                          10 New capture
```

Figure 5–25. Specifying a pattern to search for.

Forcing the Interpretation of an Embedded Protocol in an Unexpected Location

On occasion you may encounter frames that contain a known protocol embedded in such a way that the interpreters don't recognize it. This may happen when two applications exchange data in a nonstandard way, or in a way that hasn't yet become sufficiently standard for the protocol interpreter suites to know about it.

When that happens, the analyzer may simply report that (in some lower-level protocol) it has found "*xx* bytes of data." But you have reason to believe that, somewhere in that data field, there is an embedded protocol. Moreover, it's a protocol that the analyzer could perfectly well interpret— if only it knew it was there.

The **Force Protocol** feature gives you a way to tell the analyzer where to look for a protocol that it does not automatically recognize. **Force Protocol** works:

- Only after you first enable the **Force Protocol** feature.
- Only from within certain protocols that permit protocol forcing.
- Only when, working from the **Summary** view, you identify

— The outer protocol (the one that contains an unexpected[1] embedded protocol)

— The embedded protocol (and where it's located)

— Which frames should be interpreted this way.

To force interpretation of an embedded protocol

1. Capture some frames that you believe contain a known embedded protocol in a location that the protocol interpreters don't expect.

2. In the **Summary** branch of the **Display** menu, enable protocol forcing. Move the highlight to **Force Protocol** and press Spacebar to toggle between X (inactive) and ✓ (active).

   <u>Result</u>: The analyzer enables a special use of the F3 key that operates only within a **Summary** display. (Usually, F3 starts display. Now, once display has started, you'll have a new temporary meaning of F3.) The special use of F3 is labelled **Force Protocol**. This meaning of F3 is in effect (and its label shows) when:

   — The **Summary** view is open and is the active view

   and

   — The highlighted line within the **Summary** view contains a protocol from which forcing is permitted.

3. If you haven't yet started the display, press F3 to display frames in the capture buffer. To make use of **Force Protocol,** you *must* open the **Summary** view, and it's generally useful to open **Detail** and **Hex** as well.

4. Scroll until you find a frame that contains the embedded protocol.

   Depending on the circumstances, you may have various clues to identify the frame. You may be able to recognize some of its embedded data in the **Hex** view. You may recognize a particular socket as a destination in the **Summary** or the **Detail** view. The application may adopt a suggestive NetBIOS name as the destination of traffic containing the protocol. You may have to do some experimentation.

[1] Of course, what is "unexpected" depends on point of view. *You* may expect the protocol to be just where it is. But the protocol interpreter suite you're using doesn't.

5. In the **Summary** view, place the highlight on the line with the protocol whose data field contains the embedded protocol. Provided the highlight is on a protocol from which forcing is possible, you'll see that `F3 Force Protocl` is visible at the bottom of the screen (Figure 5–26).

   (Almost always, the highest level that the analyzer can interpret is the one that contains the unexpected protocol. From that point of view, it would be sufficient to use the **Summary** view with **Highest level only**. However, once you force an additional level of interpretation, you probably want to see both the new protocol and those that contain it. For this reason, it is usually preferable to deactivate **Highest level only**.)

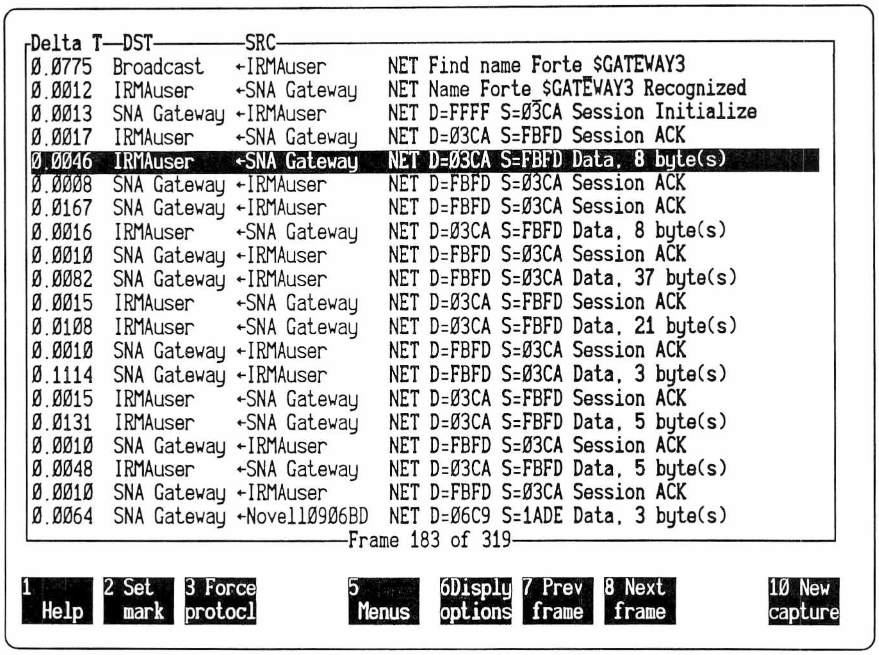

```
Delta T—DST————————SRC————————————————————————————————————————————
Ø.Ø775  Broadcast   ←IRMAuser      NET Find name Forte_$GATEWAY3
Ø.ØØ12  IRMAuser    ←SNA Gateway   NET Name Forte_$GATEWAY3 Recognized
Ø.ØØ13  SNA Gateway ←IRMAuser      NET D=FFFF S=Ø3CA Session Initialize
Ø.ØØ17  IRMAuser    ←SNA Gateway   NET D=Ø3CA S=FBFD Session ACK
Ø.ØØ46  IRMAuser    ←SNA Gateway   NET D=Ø3CA S=FBFD Data, 8 byte(s)
Ø.ØØØ8  SNA Gateway ←IRMAuser      NET D=FBFD S=Ø3CA Session ACK
Ø.Ø167  SNA Gateway ←IRMAuser      NET D=FBFD S=Ø3CA Session ACK
Ø.ØØ16  IRMAuser    ←SNA Gateway   NET D=Ø3CA S=FBFD Data, 8 byte(s)
Ø.ØØ1Ø  SNA Gateway ←IRMAuser      NET D=FBFD S=Ø3CA Session ACK
Ø.ØØ82  SNA Gateway ←IRMAuser      NET D=FBFD S=Ø3CA Data, 37 byte(s)
Ø.ØØ15  IRMAuser    ←SNA Gateway   NET D=Ø3CA S=FBFD Session ACK
Ø.Ø1Ø8  IRMAuser    ←SNA Gateway   NET D=Ø3CA S=FBFD Data, 21 byte(s)
Ø.ØØ1Ø  SNA Gateway ←IRMAuser      NET D=FBFD S=Ø3CA Session ACK
Ø.1114  SNA Gateway ←IRMAuser      NET D=FBFD S=Ø3CA Data, 3 byte(s)
Ø.ØØ15  IRMAuser    ←SNA Gateway   NET D=Ø3CA S=FBFD Session ACK
Ø.Ø131  IRMAuser    ←SNA Gateway   NET D=Ø3CA S=FBFD Data, 5 byte(s)
Ø.ØØ1Ø  SNA Gateway ←IRMAuser      NET D=FBFD S=Ø3CA Session ACK
Ø.ØØ48  IRMAuser    ←SNA Gateway   NET D=Ø3CA S=FBFD Data, 5 byte(s)
Ø.ØØ1Ø  SNA Gateway ←IRMAuser      NET D=FBFD S=Ø3CA Session ACK
Ø.ØØ64  SNA Gateway ←Novell Ø9Ø6BD NET D=Ø6C9 S=1ADE Data, 3 byte(s)
——————————————————————————————Frame 183 of 319————————————————————

1      2 Set  3 Force          5      6Disply 7 Prev  8 Next          1Ø New
 Help   mark  protocl          Menus  options  frame   frame          capture
```

Figure 5–26. **Summary** *view with F3 =* **Force Protocol** *enabled.*

6. Press F3, **Force Protocol**.

   Result: The analyzer puts up the **Force Protocol** menu, temporarily overlaying the regular display (Figure 5–27).

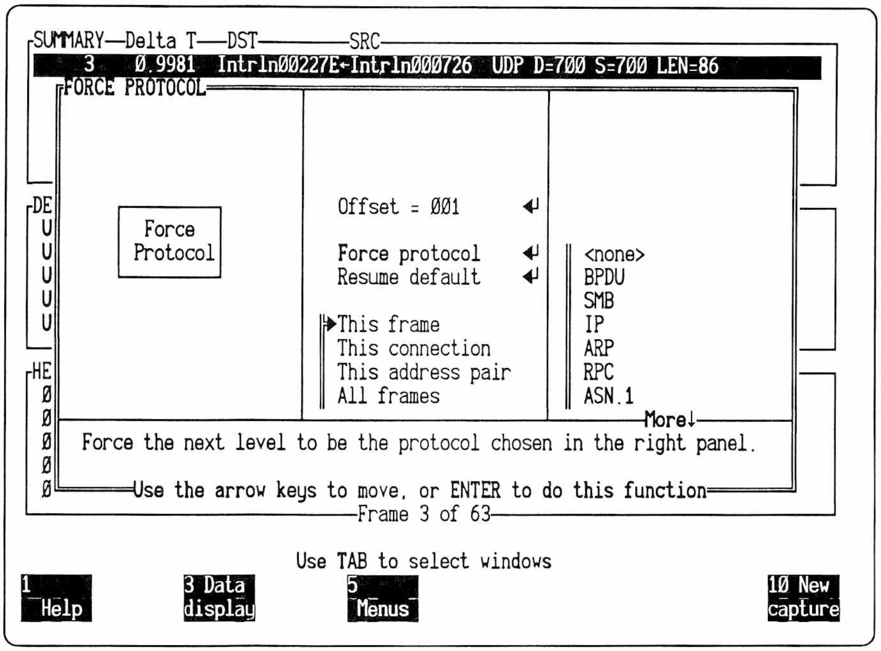

```
SUMMARY—Delta T——DST———————SRC——————————————————————————————
   3    0.9981  Intrln00227E←Intrln000726  UDP D=700 S=700 LEN=86
 FORCE PROTOCOL
                             Offset = 001        ↵
      Force
      Protocol               Force protocol      ↵   <none>
                             Resume default      ↵   BPDU
                                                     SMB
                            ▸This frame              IP
                             This connection         ARP
                             This address pair       RPC
                             All frames              ASN.1
                                                         More↓
   Force the next level to be the protocol chosen in the right panel.

        Use the arrow keys to move, or ENTER to do this function
                                Frame 3 of 63

                         Use TAB to select windows
1          3 Data      5                                    10 New
  Help     display      Menus                               capture
```

Figure 5–27. Force Protocol menu overlaid on the Display view

7. Describe the protocol you want the analyzer to use.

   a. In the panel to right, select the protocol you want the analyzer to use.

      Move the highlight to the one you want and press Spacebar. The list of alternatives includes all the protocols available through the set of protocol interpreter suites installed in this analyzer.[1]

   b. Set the offset at which the protocol starts. Move the highlight to **Offset=** and press Enter. The analyzer opens a dialog box in which you can enter the amount of the offset. When you've made your entry, press Enter to record it.

      The offset is the distance from the start of the data field in the outer protocol; that is, the distance from the end of the outer protocol's header. (The outer protocol is the one you highlighted in the **Summary** view.) Write the offset in hexadecimal.

8. Tell the analyzer which frames to interpret using the forced protocol. Move the highlight to the one you want and press Spacebar to select one of the choices. These are the four alternatives (also visible in Figure 5–27):

[1] If a protocol you want to use isn't on the list, you can use the configuration utility to build an interpreter with the right combination; see Chapter 7.

This frame
The protocol will be forced for the currently highlighted frame, but not for any others.

This connection
The protocol will be forced for the currently highlighted frame, and for all frames earlier or later in the capture buffer that are part of the same connection or logical call.

This address pair
The protocol will be forced for all frames in the capture buffer passing in either direction between the DLC source and destination of the currently highlighted frame

All frames
The protocol will be forced for all frames now in the capture buffer.

9. Apply the protocol you've specified— or abandon it.

   To return to the display, with the protocol forced as you've specified, move the highlight to **Force protocol** and press Enter.

   To return to the display without forcing (but leaving the protocol and offset as you've selected), press Esc.

10. Inspect the display for confirmation that you have forced the protocol interpretation correctly.

    Some confirming signs:

    — The forced protocol makes sense, revealing reasonable commands, addresses, and data.

    — Protocols within the previously-hidden protocol become visible and make sense.

    Some disconfirming signs:

    — The detail view of the forced protocol shows an invalid checksum (in a protocol that provides its own checksum).

    — The embedded protocol is shorter than the data field that contains it— not conclusive, but suspicious.

    — The embedded protocol is too short. The last line of the protocol's **Detail** display will indicate when the interpreter reached the end of the data before the normal end for this protocol.

Undoing a Line of Forced Protocol

Protocol forcing has a temporary effect on the display. If you close the display, or replace the contents of the capture buffer, the analyzer discards any protocol forcing you did earlier. When you save the capture buffer, the saved file does not contain any record of the protocol forcing you may have applied to it.

While a display is active, the **Force Protocol** menu gives you a way to undo the forcing of individual frames.

To undo a forced protocol

1. Display the capture buffer using the **Summary** view with **Highest level only** turned off. (You can have other views open or not, as you wish.)
2. Move the highlight to the line showing the outer protocol—that is, the protocol from which you forced interpretation.

   (For example, if you earlier highlighted NetBIOS and forced the interpretation of SNA, you should now highlight the NetBIOS line, not the SNA now visible below it.)
3. Press F3, **Force Protocol**.

   Result: The analyzer overlays the **Force Protocol** menu.
4. Move the highlight to **Resume default**, and press Enter.

   Result: The analyzer closes the **Force protocol** overlay and returns to the display. The interpretation of the highlighted line returns to its default.

Forced Protocol Example: SNA Embedded within NetBIOS

A personal computer on an Ethernet network uses an emulator that makes it appear as a 327x terminal to the 3274 controller on an IBM mainframe. The PC's emulator card uses Novell NetBIOS over XNS. Once it has set up a NetBIOS connection to the controller, the terminal emulator embeds SNA commands in IRMA protocol, and transmits these as the data field of a NetBIOS packet. The interpreter knows how to recognize SNA within an IRMA packet, but it doesn't know how to tell that a NetBIOS packet will contain IRMA. The presence of an IRMA packet is a convention of the two ends of the connection. There is no header that announces "The rest of this field is IRMA data."

To decode the encapsulated IRMA and SNA data, you might proceed as follows:

1. Capture frames from the PC using the terminal emulator. Enable Force Protocol. Display the captured frames with the standard protocol interpreters, including Novell XNS and NetBIOS.

   The display shows NetBIOS transmission between the PC and the SNA gateway, but doesn't decode them (Figure 5–28).

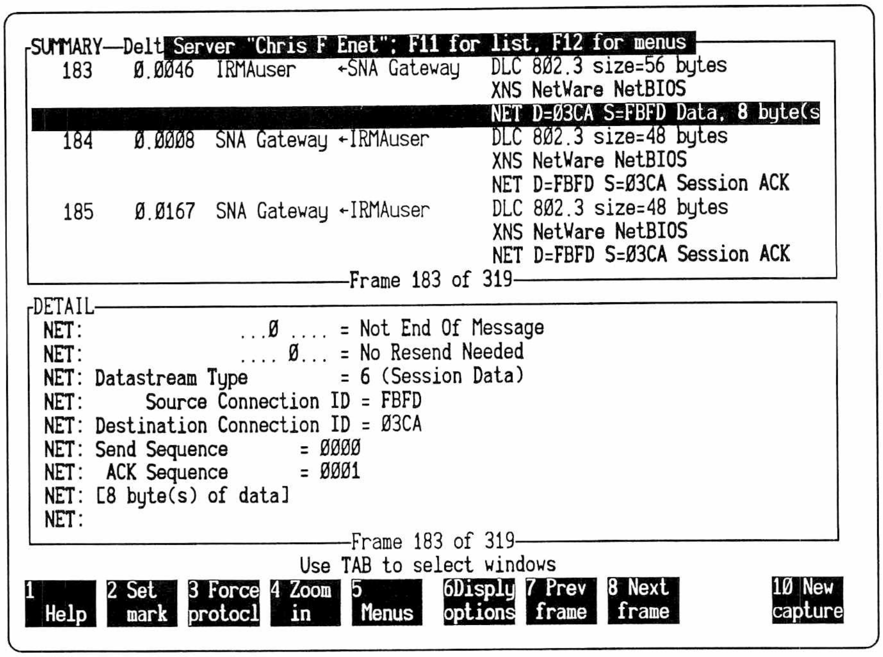

```
SUMMARY—Delt Server "Chris F Enet": F11 for list, F12 for menus
 183    Ø.ØØ46  IRMAuser    ←SNA Gateway    DLC 8Ø2.3 size=56 bytes
                                            XNS NetWare NetBIOS
                                            NET D=Ø3CA S=FBFD Data, 8 byte(s
 184    Ø.ØØØ8  SNA Gateway ←IRMAuser       DLC 8Ø2.3 size=48 bytes
                                            XNS NetWare NetBIOS
                                            NET D=FBFD S=Ø3CA Session ACK
 185    Ø.Ø167  SNA Gateway ←IRMAuser       DLC 8Ø2.3 size=48 bytes
                                            XNS NetWare NetBIOS
                                            NET D=FBFD S=Ø3CA Session ACK
                              Frame 183 of 319
DETAIL
 NET:                    ...Ø .... = Not End Of Message
 NET:                    .... Ø... = No Resend Needed
 NET: Datastream Type              = 6 (Session Data)
 NET:     Source Connection ID = FBFD
 NET: Destination Connection ID = Ø3CA
 NET: Send Sequence           = ØØØØ
 NET:  ACK Sequence           = ØØØ1
 NET: [8 byte(s) of data]
 NET:
                              Frame 183 of 319
                          Use TAB to select windows
```

Figure 5–28. Uninterpreted NetBIOS data from a terminal emulator.

2. Identify a frame that probably contains SNA data. (In this case, there are two clues. First, the PC establishes a connection with a NetBIOS address named "Forte Gateway". Second, by forcing EBCDIC rather than ASCII display in the Hex view, some SNA messages are apparent in the NetBIOS data field.

3. With the highlight on the line produced by the NetBIOS interpreter (labelled "NET" in the left margin), press F3 **Force Protocol.**[1]

4. In the **Force Protocol** menu, select IRMA protocol at offset 0, and apply it to this frame only. (You could select SNA at offset 3, and thereby jump over the three bytes of IRMA.) If that seems to work, you can then select **This connection,** since the protocol probably applies throughout this NetBIOS connection.

[1] The summary of the preceding line says "NetWare NetBIOS." That's the interpretation of the NetWare frame saying that "NetBIOS follows", and not the NetBIOS frame itself.

The result is visible in Figure 5–29. The IRMA protocol becomes visible as a result of forcing, and the SNA appears automatically since it is embedded in the IRMA frame in a standard manner.

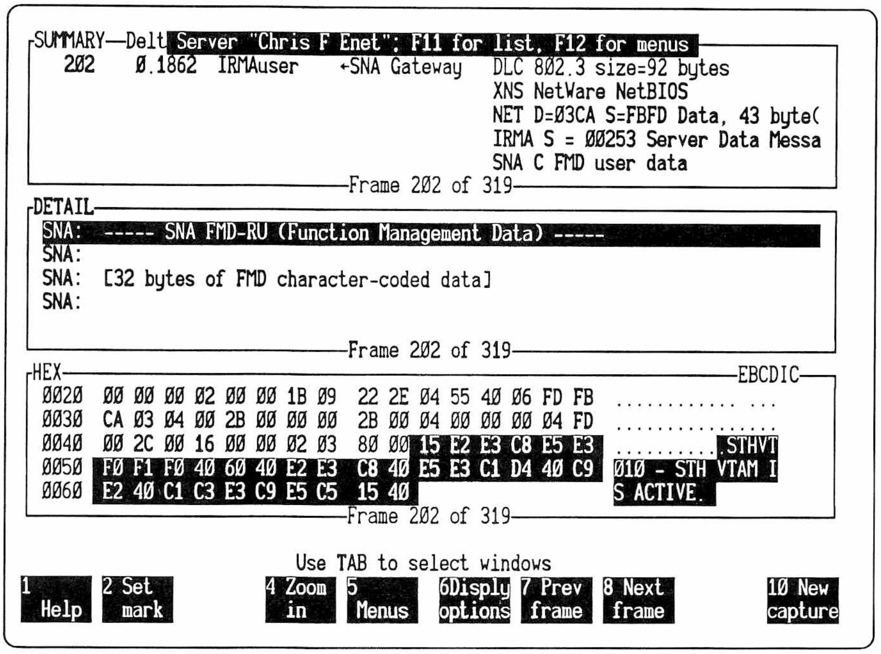

```
SUMMARY—Delt Server "Chris F Enet": F11 for list, F12 for menus
  202     Ø.1862  IRMAuser     ←SNA Gateway   DLC 8Ø2.3 size=92 bytes
                                              XNS NetWare NetBIOS
                                              NET D=Ø3CA S=FBFD Data, 43 byte(
                                              IRMA S = ØØ253 Server Data Messa
                                              SNA C FMD user data
                              Frame 2Ø2 of 319
DETAIL
 SNA:  ----- SNA FMD-RU (Function Management Data) -----
 SNA:
 SNA:  [32 bytes of FMD character-coded data]
 SNA:
                              Frame 2Ø2 of 319
HEX                                                              EBCDIC
 ØØ2Ø  ØØ ØØ ØØ Ø2 ØØ ØØ 1B Ø9  22 2E Ø4 55 4Ø Ø6 FD FB  .............  ...
 ØØ3Ø  CA Ø3 Ø4 ØØ 2B ØØ ØØ ØØ  2B ØØ Ø4 ØØ ØØ ØØ Ø4 FD  ................
 ØØ4Ø  ØØ 2C ØØ 16 ØØ ØØ Ø2 Ø3  8Ø ØØ 15 E2 E3 C8 E5 E3  ...........STHVT
 ØØ5Ø  FØ F1 FØ 4Ø 6Ø 4Ø E2 E3  C8 4Ø E5 E3 C1 D4 4Ø C9  Ø1Ø - STH VTAM I
 ØØ6Ø  E2 4Ø C1 C3 E3 C9 E5 C5  15 4Ø                    S ACTIVE.
                              Frame 2Ø2 of 319
```

Figure 5–29. Forced protocol reveals IRMA and SNA within NetBIOS.

Exporting a Report on Frames in the Capture Buffer

You can generate a report on the frames in the capture buffer. The report can be sent directly to the printer, or to a file. The file can be formatted for printing or for input to another program (for example, to a spread sheet).

Frames that pass your current display filters are included in the report. The fields reported are those currently displayed on the screen. However, since the report isn't confined by the size of a screen panel, it doesn't exactly duplicate the layout you'd see on screen. For example, reports do not attempt to represent colors or highlighting.

Range of Frames Included in the Report

The report includes all displayable frames within a certain range. By default, the Sniffer analyzer assumes the range includes all displayable frames, from the first to the last.

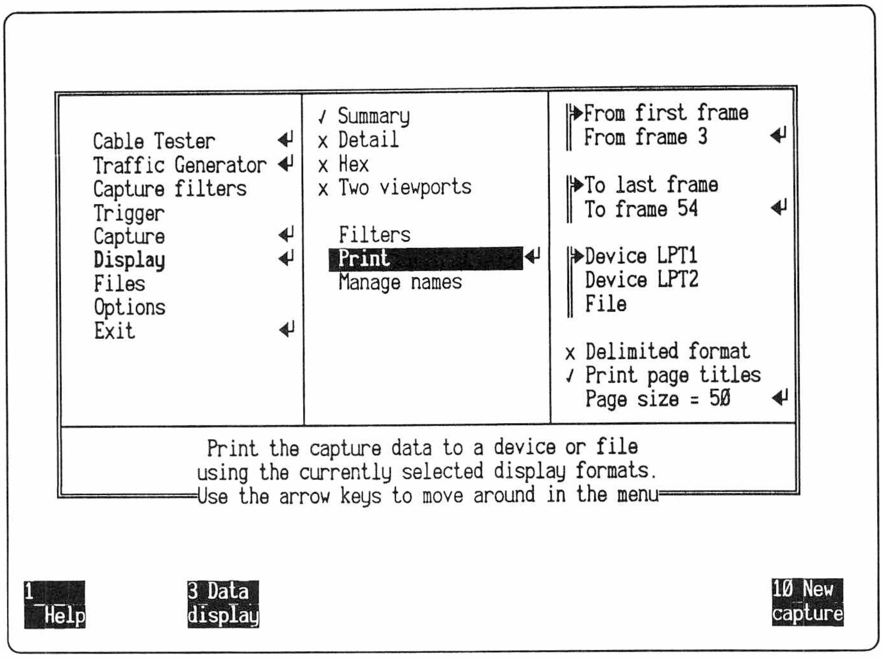

```
 Cable Tester       ↵ | ✓ Summary          | ▶From first frame
 Traffic Generator  ↵ | x Detail           |  From frame 3       ↵
 Capture filters      | x Hex              |
 Trigger              | x Two viewports    | ▶To last frame
 Capture            ↵ |                    |  To frame 54        ↵
 Display            ↵ |   Filters          |
 Files                |   Print          ↵ | ▶Device LPT1
 Options              |   Manage names     |  Device LPT2
 Exit               ↵ |                    |  File
                      |                    |
                      |                    | x Delimited format
                      |                    | ✓ Print page titles
                      |                    |   Page size = 50    ↵

              Print the capture data to a device or file
             using the currently selected display formats.
            Use the arrow keys to move around in the menu

 1            3 Data                                       10 New
   Help       display                                      capture
```

Figure 5–30. Options for printing to device or file.

Alternatively, you can limit the range by setting the frame numbers of the first and last frames to be printed. A pair of radio controls flips between **First** or **Last** and a specific number.

The Sniffer analyzer initializes the display with its guess about the frame numbers you might want. It assumes that (if you choose something other than first to last) you'd want frames lying between the *marked* frame and the *current* frame (Figure 5–30). (You mark a frame by pressing F2 while it is highlighted; if you haven't set the mark, the mark is on frame 1. The *current* frame is the one now highlighted in the display.)

To change the printing range, highlight the phrase **From frame** or **To frame** and press Enter. The analyzer opens a dialog box for you to supply a frame number.

Where the Output Goes: Printer or File

You have three possible destinations for the printer output:

LPT1 A printer directly attached to the Sniffer server's LPT1.

LPT2 A printer directly attached to the SniffMaster console from which you are operating the server. (The console may then redirect the output to a file on its own hard disk, if it has been configured to do so.) For a discussion of console configurations, see *Distributed Sniffer System: Installation and Operations Manual.*

File A file that you name on the Sniffer server's hard disk.

For more details of the physical setup of attached printers, see *Distributed Sniffer System: Installation and Operations Manual.*

Figure 5–31 shows a printed report from an Ethernet network. This report shows part of the first page of a report on the **Summary** view in two-station format.[1]

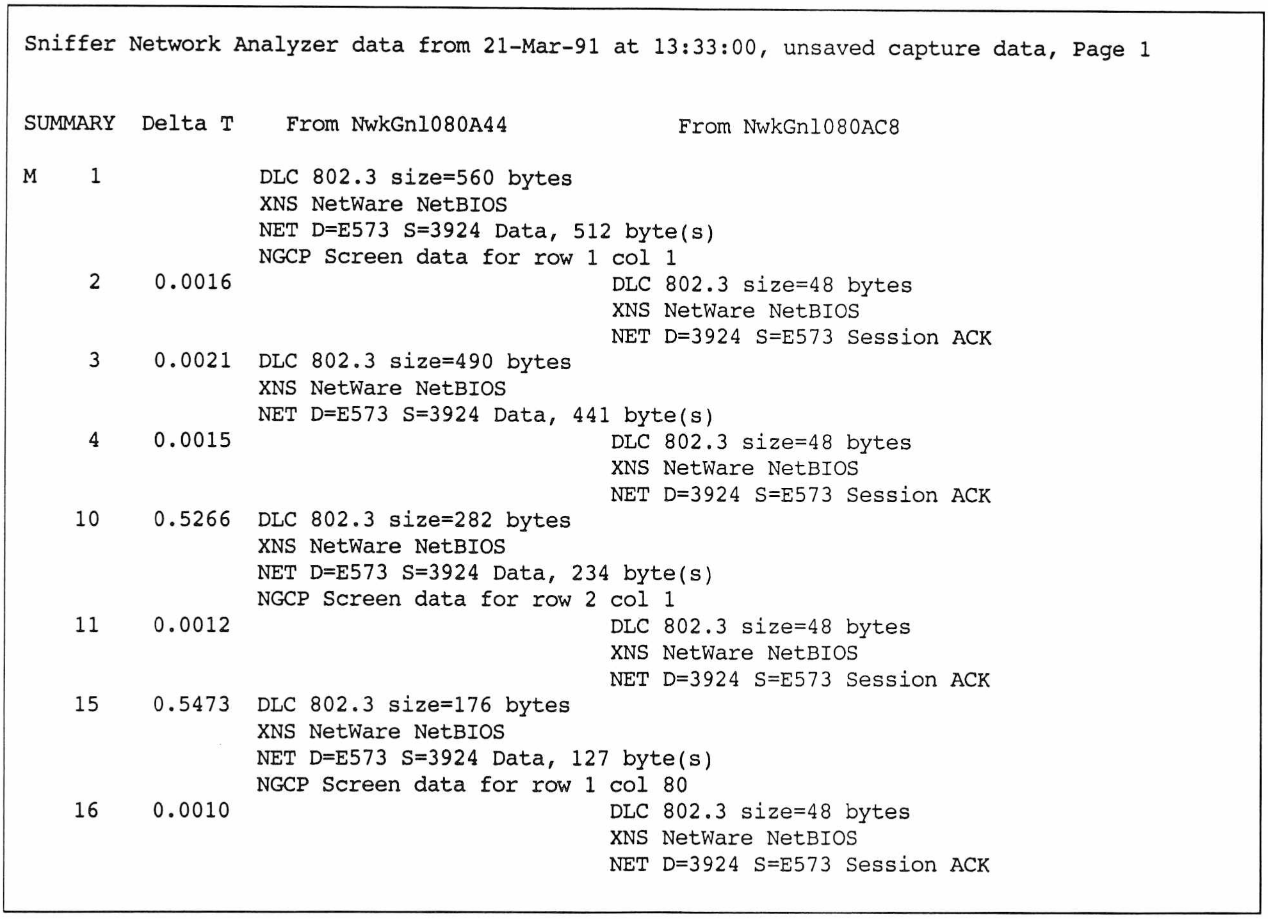

```
Sniffer Network Analyzer data from 21-Mar-91 at 13:33:00, unsaved capture data, Page 1


SUMMARY  Delta T    From NwkGnl080A44            From NwkGnl080AC8

M    1            DLC 802.3 size=560 bytes
                  XNS NetWare NetBIOS
                  NET D=E573 S=3924 Data, 512 byte(s)
                  NGCP Screen data for row 1 col 1
     2   0.0016                               DLC 802.3 size=48 bytes
                                              XNS NetWare NetBIOS
                                              NET D=3924 S=E573 Session ACK
     3   0.0021   DLC 802.3 size=490 bytes
                  XNS NetWare NetBIOS
                  NET D=E573 S=3924 Data, 441 byte(s)
     4   0.0015                               DLC 802.3 size=48 bytes
                                              XNS NetWare NetBIOS
                                              NET D=3924 S=E573 Session ACK
    10   0.5266   DLC 802.3 size=282 bytes
                  XNS NetWare NetBIOS
                  NET D=E573 S=3924 Data, 234 byte(s)
                  NGCP Screen data for row 2 col 1
    11   0.0012                               DLC 802.3 size=48 bytes
                                              XNS NetWare NetBIOS
                                              NET D=3924 S=E573 Session ACK
    15   0.5473   DLC 802.3 size=176 bytes
                  XNS NetWare NetBIOS
                  NET D=E573 S=3924 Data, 127 byte(s)
                  NGCP Screen data for row 1 col 80
    16   0.0010                               DLC 802.3 size=48 bytes
                                              XNS NetWare NetBIOS
                                              NET D=3924 S=E573 Session ACK
```

Figure 5–31. Portion of printed report of a **Summary** *view.*

Print to File

You can direct the printed output to a file on the Sniffer server's hard disk. (If you subsequently want to transfer the file from the server to the console, use the file transfer utility, described in *Distributed Sniffer System: Installation and Operations Manual.*)

[1] The example shows part of the control exchange between a Sniffer analysis server and a SniffMaster console, using NGCP protocol.

Content of the Exported Files

In general, a report sent to the printer or to a file contains the same information you see on the screen, but without the restriction of the small panel. **Detail** sent to a file is treated differently from **Detail** on screen. On screen, only a small amount of **Detail** is visible, but you can scroll to any level. In an exported file, scrolling is not an issue. However, in a file, **Detail** includes only the protocols you request. See Figure 5–32.

Displayed on screen	Sent to file or printer
Summary shows a line for each protocol selected in the protocol or address level filters.	
When a frame's **Summary** is highlighted in the upper panel, its **Detail** is visible in the lower panel.	A frame's entire **Summary** appears first, and then its **Detail**.
A frame's **Detail** starts with the level that is highlighted in the **Summary** view. You can scroll to see the interpretation of any protocol the frame contains, regardless of display filters.	**Detail** is limited to the protocols selected in the **Protocols** filter. **Detail** follows the **Summary** for each frame, with protocols in order from lowest to highest.

Figure 5–32. Treatment of **Summary** *and* **Detail** *on screen and to a file.*

Delimited Format for Export to a Spread Sheet

When you select **Delimited format**, the output is written as *comma-separated values*. In this format, each character field is surrounded by double quotes. A numeric field is written as ASCII characters without quotes. Successive fields are separated by commas.

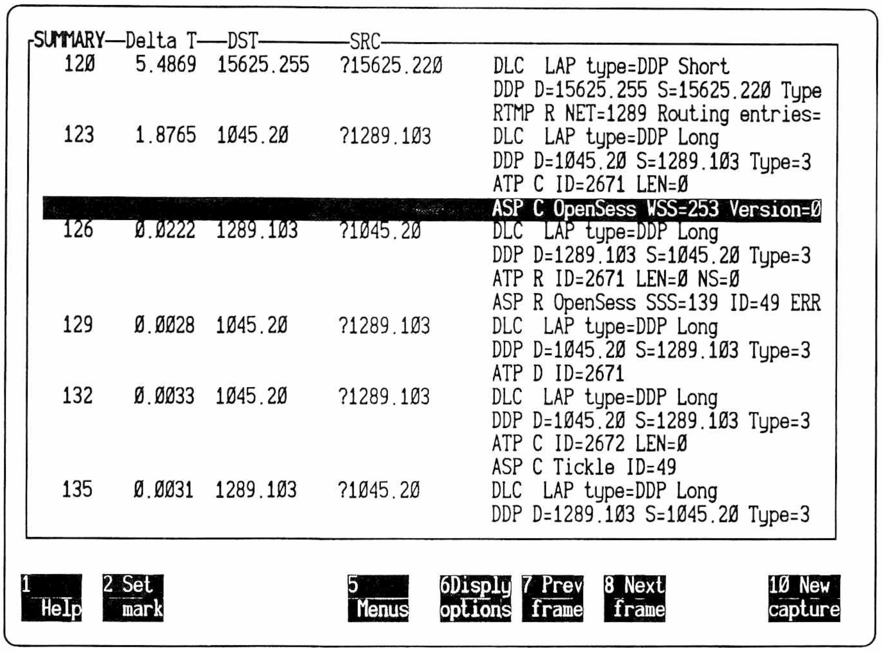

```
SUMMARY—Delta T——DST—————SRC———————————————————————————————————
  12Ø    5.4869  15625.255   ?15625.22Ø    DLC  LAP type=DDP Short
                                           DDP D=15625.255 S=15625.22Ø Type
                                           RTMP R NET=1289 Routing entries=
  123    1.8765  1Ø45.2Ø     ?1289.1Ø3     DLC  LAP type=DDP Long
                                           DDP D=1Ø45.2Ø S=1289.1Ø3 Type=3
                                           ATP C ID=2671 LEN=Ø
                                           ASP C OpenSess WSS=253 Version=Ø
  126    Ø.Ø222  1289.1Ø3    ?1Ø45.2Ø      DLC  LAP type=DDP Long
                                           DDP D=1289.1Ø3 S=1Ø45.2Ø Type=3
                                           ATP R ID=2671 LEN=Ø NS=Ø
                                           ASP R OpenSess SSS=139 ID=49 ERR
  129    Ø.ØØ28  1Ø45.2Ø     ?1289.1Ø3     DLC  LAP type=DDP Long
                                           DDP D=1Ø45.2Ø S=1289.1Ø3 Type=3
                                           ATP D ID=2671
  132    Ø.ØØ33  1Ø45.2Ø     ?1289.1Ø3     DLC  LAP type=DDP Long
                                           DDP D=1Ø45.2Ø S=1289.1Ø3 Type=3
                                           ATP C ID=2672 LEN=Ø
                                           ASP C Tickle ID=49
  135    Ø.ØØ31  1289.1Ø3    ?1Ø45.2Ø      DLC  LAP type=DDP Long
                                           DDP D=1289.1Ø3 S=1Ø45.2Ø Type=3
```

Figure 5–33. Portion of a **Summary** *display on screen. Compare with CSV format of the same data in Figure 5–34.*

Comma-separated values is a format widely used for importing data to spread sheets. The file's first line defines the fields. Each subsequent line is a *record,* containing the value of each field. A line of the exported file corresponds to a row of the **Summary** view on screen. That is, the file has either one line per frame or one line per protocol level.

When you turn off highest level only, a frame's display may take several rows. On screen, a field that is the same for every row (for example, the frame number) is written only once, on the first of the frame's rows. That field is left blank in the frame's later rows. However, in the comma-separated values file, every row is filled in, even when a field repeats the row above. Figure 5–33 and Figure 5–34 show screen and CSV versions of the same data.

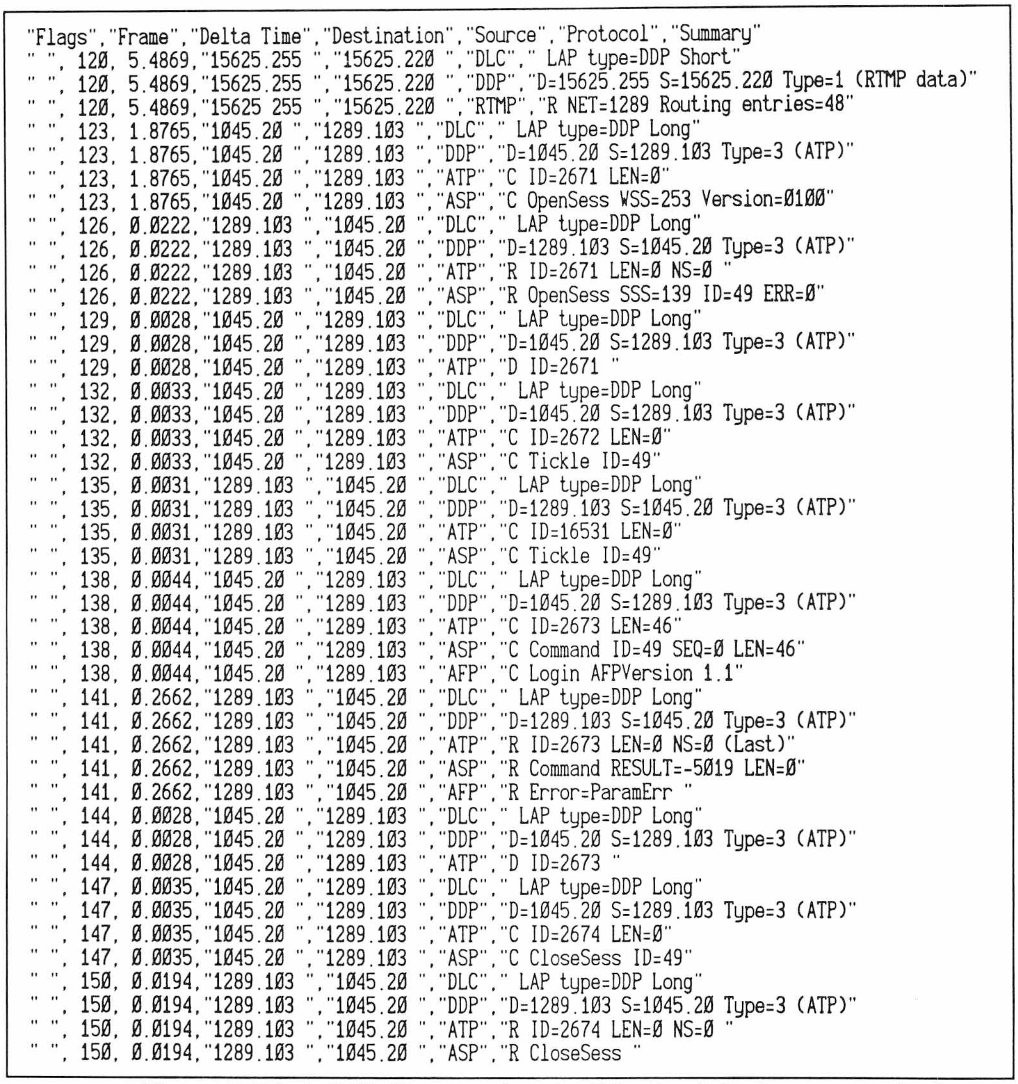

```
"Flags","Frame","Delta Time","Destination","Source","Protocol","Summary"
" ", 12Ø, 5.4869,"15625.255 ","15625.22Ø ","DLC"," LAP type=DDP Short"
" ", 12Ø, 5.4869,"15625.255 ","15625.22Ø ","DDP","D=15625.255 S=15625.22Ø Type=1 (RTMP data)"
" ", 12Ø, 5.4869,"15625 255 ","15625.22Ø ","RTMP","R NET=1289 Routing entries=48"
" ", 123, 1.8765,"1Ø45.2Ø ","1289.1Ø3 ","DLC"," LAP type=DDP Long"
" ", 123, 1.8765,"1Ø45.2Ø ","1289.1Ø3 ","DDP","D=1Ø45.2Ø S=1289.1Ø3 Type=3 (ATP)"
" ", 123, 1.8765,"1Ø45.2Ø ","1289.1Ø3 ","ATP","C ID=2671 LEN=Ø"
" ", 123, 1.8765,"1Ø45.2Ø ","1289.1Ø3 ","ASP","C OpenSess WSS=253 Version=Ø1ØØ"
" ", 126, Ø.Ø222,"1289.1Ø3 ","1Ø45.2Ø ","DLC"," LAP type=DDP Long"
" ", 126, Ø.Ø222,"1289.1Ø3 ","1Ø45.2Ø ","DDP","D=1289.1Ø3 S=1Ø45.2Ø Type=3 (ATP)"
" ", 126, Ø.Ø222,"1289.1Ø3 ","1Ø45.2Ø ","ATP","R ID=2671 LEN=Ø NS=Ø "
" ", 126, Ø.Ø222,"1289.1Ø3 ","1Ø45.2Ø ","ASP","R OpenSess SSS=139 ID=49 ERR=Ø"
" ", 129, Ø.ØØ28,"1Ø45.2Ø ","1289.1Ø3 ","DLC"," LAP type=DDP Long"
" ", 129, Ø.ØØ28,"1Ø45.2Ø ","1289.1Ø3 ","DDP","D=1Ø45.2Ø S=1289.1Ø3 Type=3 (ATP)"
" ", 129, Ø.ØØ28,"1Ø45.2Ø ","1289.1Ø3 ","ATP","D ID=2671 "
" ", 132, Ø.ØØ33,"1Ø45.2Ø ","1289.1Ø3 ","DLC"," LAP type=DDP Long"
" ", 132, Ø.ØØ33,"1Ø45.2Ø ","1289.1Ø3 ","DDP","D=1Ø45.2Ø S=1289.1Ø3 Type=3 (ATP)"
" ", 132, Ø.ØØ33,"1Ø45.2Ø ","1289.1Ø3 ","ATP","C ID=2672 LEN=Ø"
" ", 132, Ø.ØØ33,"1Ø45.2Ø ","1289.1Ø3 ","ASP","C Tickle ID=49"
" ", 135, Ø.ØØ31,"1289.1Ø3 ","1Ø45.2Ø ","DLC"," LAP type=DDP Long"
" ", 135, Ø.ØØ31,"1289.1Ø3 ","1Ø45.2Ø ","DDP","D=1289.1Ø3 S=1Ø45.2Ø Type=3 (ATP)"
" ", 135, Ø.ØØ31,"1289.1Ø3 ","1Ø45.2Ø ","ATP","C ID=16531 LEN=Ø"
" ", 135, Ø.ØØ31,"1289.1Ø3 ","1Ø45.2Ø ","ASP","C Tickle ID=49"
" ", 138, Ø.ØØ44,"1Ø45.2Ø ","1289.1Ø3 ","DLC"," LAP type=DDP Long"
" ", 138, Ø.ØØ44,"1Ø45.2Ø ","1289.1Ø3 ","DDP","D=1Ø45.2Ø S=1289.1Ø3 Type=3 (ATP)"
" ", 138, Ø.ØØ44,"1Ø45.2Ø ","1289.1Ø3 ","ATP","C ID=2673 LEN=46"
" ", 138, Ø.ØØ44,"1Ø45.2Ø ","1289.1Ø3 ","ASP","C Command ID=49 SEQ=Ø LEN=46"
" ", 138, Ø.ØØ44,"1Ø45.2Ø ","1289.1Ø3 ","AFP","C Login AFPVersion 1.1"
" ", 141, Ø.2662,"1289.1Ø3 ","1Ø45.2Ø ","DLC"," LAP type=DDP Long"
" ", 141, Ø.2662,"1289.1Ø3 ","1Ø45.2Ø ","DDP","D=1289.1Ø3 S=1Ø45.2Ø Type=3 (ATP)"
" ", 141, Ø.2662,"1289.1Ø3 ","1Ø45.2Ø ","ATP","R ID=2673 LEN=Ø NS=Ø (Last)"
" ", 141, Ø.2662,"1289.1Ø3 ","1Ø45.2Ø ","ASP","R Command RESULT=-5Ø19 LEN=Ø"
" ", 141, Ø.2662,"1289.1Ø3 ","1Ø45.2Ø ","AFP","R Error=ParamErr "
" ", 144, Ø.ØØ28,"1Ø45.2Ø ","1289.1Ø3 ","DLC"," LAP type=DDP Long"
" ", 144, Ø.ØØ28,"1Ø45.2Ø ","1289.1Ø3 ","DDP","D=1Ø45.2Ø S=1289.1Ø3 Type=3 (ATP)"
" ", 144, Ø.ØØ28,"1Ø45.2Ø ","1289.1Ø3 ","ATP","D ID=2673 "
" ", 147, Ø.ØØ35,"1Ø45.2Ø ","1289.1Ø3 ","DLC"," LAP type=DDP Long"
" ", 147, Ø.ØØ35,"1Ø45.2Ø ","1289.1Ø3 ","DDP","D=1Ø45.2Ø S=1289.1Ø3 Type=3 (ATP)"
" ", 147, Ø.ØØ35,"1Ø45.2Ø ","1289.1Ø3 ","ATP","C ID=2674 LEN=Ø"
" ", 147, Ø.ØØ35,"1Ø45.2Ø ","1289.1Ø3 ","ASP","C CloseSess ID=49"
" ", 15Ø, Ø.Ø194,"1289.1Ø3 ","1Ø45.2Ø ","DLC"," LAP type=DDP Long"
" ", 15Ø, Ø.Ø194,"1289.1Ø3 ","1Ø45.2Ø ","DDP","D=1289.1Ø3 S=1Ø45.2Ø Type=3 (ATP)"
" ", 15Ø, Ø.Ø194,"1289.1Ø3 ","1Ø45.2Ø ","ATP","R ID=2674 LEN=Ø NS=Ø "
" ", 15Ø, Ø.Ø194,"1289.1Ø3 ","1Ø45.2Ø ","ASP","R CloseSess "
```

Figure 5–34. A CSV file created with all levels of the **Summary** *display.*

(It isn't illegal to ask for CSV format while reporting **Detail** or **Hex** views. However, the output is in CSV format only for the **Summary** view. The CSV file includes data from the other views, but in the standard printer format.)

Page Titles

If you check the page titles option, each page (whether written to a file or directly to the printer) starts with a heading. The heading specifies the date and time the data was recorded. It also shows the name of the network (when the file was captured live) or the name of the file (when the capture buffer was loaded from a saved file). At the far right, it shows the page number.

There are two blank lines between the heading and the start of the data.

Having page titles also causes explicit page breaks. When you select page titles, the Sniffer analyzer includes a form feed character after the last non-blank line of each page.

Selecting CSV format turns off page titles. (You can turn them on again if you really want; however, most applications that accept CSV format won't want page titles.)

Page Size

When you select page titles, you can also set page size, the number of lines per page. By default, the number is 50. That is the total number of printed lines. For example, if you use headings, and accept the default of 50 lines per page you get the heading, two blank lines, and 47 lines of data.

Since each page break is indicated by an explicit form feed, there is no separate setting for the physical length of the paper.

To send a report to file or printer

1. Set **Display** filters to exclude extraneous frames (page 5–5). These may include
   - Address level filter
   - Destination class filter
   - Station address filters
   - Protocol filters
   - Pattern match filters
2. Set **Address level** filters (page 5–7) to determine how source and destination will be identified in the output. The output will show each frame's source and destination at the highest of the levels you checked. (Notice that **Address level** affects the way source and destination are described even when it doesn't alter which frames are displayed.)
3. Open the views (**Summary**, **Detail** or **Hex**) that you want included in the report (page 5–14).
4. Activate or deactivate **Highest level only**, as appropriate (page 5–17).
5. Activate or deactivate **Two-station format**, as appropriate (page 5–18).
6. Set the width of the source and destination name fields (page 5–17).

The default width is 12, the minimum 8. Making these fields wider or narrower displaces the fields to the right by a corresponding amount. The analyzer does not fold the output lines (so the treatment of long lines is up to the printer or application that processes the output).

7. Set the range of frames to be printed:
   - First frame
   - Last frame

   To change the numeric frame number, move the highlight to **First frame** or **Last frame** and press Enter. The analyzer opens a dialog box to receive a new number. Type the number you want and press Enter.

8. Set the output destination:
   - LPT1 (the server's printer)
   - LPT2 (the console's printer)
   - File (on the server's hard disk)

   If you ask for output to file, when you finally move the highlight to Print and press Enter, the analyzer will prompt you to name the output file.

9. Indicate whether you want the **Summary** fields to be written in *delimited* format (page 5–54).

10. Indicate whether you want titles at the top of each page (page 5–56).

    Titles, if used, gives the date and time that data were collected, the name of the file from which data was loaded (if any), and the page number, in one line. You can't have page titles for output to file in delimited format.

11. Set the total number of printed lines on each page.

    The title line, if any, precedes the page's data, and is separated from the data by two blank lines. Data lines are thus reduced by 3 when you have titles. Each page of data is terminated by a formfeed (hex 0C).

    The default number of lines is 50. The maximum is 9999. If you set 0 lines, output is not broken into pages and there is no trailing formfeed.

12. When you've completed all the specifications (except the filename when you requested output to file) move the highlight to **Print** and press Enter.

    What happens next depends on the destination:

    LPT1 The server starts printing to the device attached to its parallel port. (If the device isn't ready, the server will — after a delay— report that fact.)

LPT2 The server starts transmitting the data back to the console, which redirects it to the console's parallel port.

File The analyzer opens a dialog box for you to write the file's name. Initially, the dialog box contains the path but no name. (To change the default path, or to create a new directory, see page 6–9.)

The path shows C (the server's hard drive) and the current directory. You can overwrite the directory if you wish. Write your filename using no more than eight characters. Don't include an extension; the Sniffer analyzer automatically attaches the extension .PRN for a printer file, or .CSV ("comma-separated values") for a file in the delimited format. If the name you propose is already in use, the analyzer warns you; you can abort the request or go ahead and overwrite the existing file.

Managing Names in the Displays

To make its displays more readable, the Sniffer analyzer permits you to substitute names for numeric addresses. The menus offer several options to supply new names, modify existing names, or preserve the use of names.

The Working Name Table

To translate between the addresses and the names that appear in the display, both the monitor and the analyzer refer to a name table. For each address, the name table contains three columns (see Figure 6–7 and Figure 6–8, page 6–15ff):

Address level The protocol in which the address can occur. The list of possible address levels depends on the protocol interpreter suites you have installed.

Station address The sequence of bytes that constitute the numeric station address.

Symbolic name The word or phrase you have assigned to the address (blank, when you haven't assigned one).

In the table, a station address never exists alone, but is always paired with a specific protocol. A symbolic name is thus an equivalent, not for an address alone, but for a particular pairing of protocol-and-address.

Dynamic Nature of the Name Table

A working name table contains two kinds of entries:

- Named protocol-and-address pairs.
- Unnamed protocol-and-address pairs (whose name field is therefore blank).

The Sniffer analyzer sorts the table alphabetically by name. The unnamed addresses, having blank names, appear first.

The working name table is built and used in a sequence of stages:

1. Initialization
2. Capture
3. First display
4. Individual frame display
5. Editing (may be done at any time)

Initialization

When it starts, the Sniffer analyzer initializes its working name table with names and addresses in the file STARTUP.xxD. (It omits addresses that lack names in the startup file.) The Sniffer analyzer also inserts its own address and assigns itself the name This Sniffer. (It does so even when the saved table previously assigned it a different name.)

There is a maximum size for the working name table. It has room for 500 names.

Capture

During capture, when you show pairwise or individual traffic counts, the Sniffer analyzer attaches a label to each counter. The label shows the station's address. When the name table already contains a name for that address, the analyzer uses the name rather than the numeric station address.

First display

The first time you use **Display** after a new capture, the Sniffer analyzer scans all the frames in the buffer for new addresses. It puts new addresses into the working name table with blank names (see page 5–65).

Display of individual frames

During display, each time the analysis server generates an entry for the **Summary** or **Detail** description of a frame, it checks whether the frame's addresses are in the name table. It adds any address it finds up to the limit of 50 unnamed addresses at each address level.

Editing

Before or after any of these stages, you can edit the working name table. You can add new addresses, insert names for addresses, or edit any of the entries.

Three Ways to Assign Names to Addresses

You manage names by editing the working name table. There are several different routes:

Edit
You can manually provide the symbolic name by editing the working name table (see page 5–62).

Resolve Name
Ask the Sniffer analyzer to search a previously-saved name file and to copy from it symbolic names for addresses that are unnamed in the working name table (see page 5–66).

Look for names
Certain protocols let stations exchange tables of names and addresses. When you ask the Sniffer analyzer to **Look for names**, it scans the capture buffer for such messages and extracts names from them (see page 5–65).

Saving the Name Table for Later Use

When you ask the Sniffer analyzer to **Save names**, it copies the current working name table to the startup file stored on its hard disk. Then each time you restart the analyzer, the working name table is primed with the names in the startup file.

When the startup file contains names that are no longer appropriate, you can edit them individually or clear them all out and start over; see the section that follows.

The changes you make affect the working name table (in the analysis server's main memory). When you exit the analysis server, it discards the working name table. (If you haven't saved the name table, the analyzer warns you that it's about to discard it. If you want to save the name table, first cancel the request to exit.) To record the names, before you exit, use the **Save names** command to copy the working table into the startup file.

Editing the Name Table

The procedure to edit the name table is simple. However, the task gets even easier when you let the analyzer compile a list of addresses for you *before* you start editing.

To edit the working name table

1. Before you open the name table, get the analyzer to compile a list of addresses that occur in your traffic. To do that:

   a. Capture some traffic (live from the network or by replaying a saved file). Press F10 to stop capture.

   b. Move the highlight to **Display**, then to **Filters**, then to **Address level**. Set the **address level** filters to activate all levels for which you want addresses

   c. Return the highlight to **Display** and press Enter. (Or press F3.) Display some of the captured frames.

      — If you're only interested in DLC addresses, it is sufficient to the display the first frame. The act of opening the display causes the analyzer to scan the entire capture buffer for DLC addresses).

      — If you're interested in higher level addresses, scroll through the capture buffer. The analyzer adds higher level addresses to the name table as it displays them.

2. Press F6 for **Display options**. Scroll down to **Manage names**. (You can also get there by F5, then **Display**, then **Manage names**.). With the highlight on **Manage names**, press Enter.

   Result: the Sniffer analyzer opens a dialog box showing the current name table (Figure 5–35). One line in the table is highlighted. You can scroll through the table to move the highlight or bring other names into view.

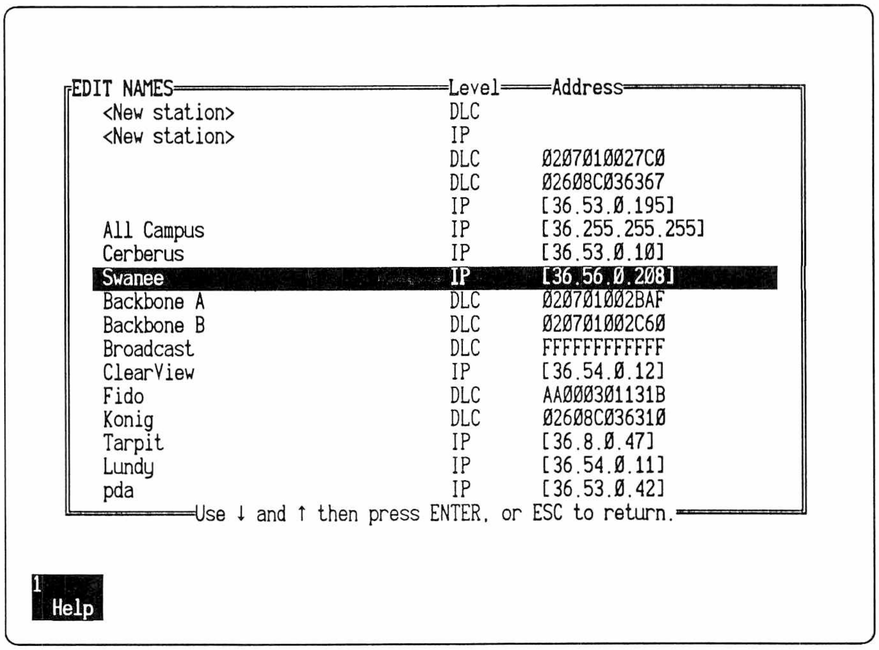

```
EDIT NAMES                              Level   Address
   <New station>                        DLC
   <New station>                        IP
                                        DLC     Ø207Ø1ØØ27CØ
                                        DLC     Ø26Ø8CØ36367
                                        IP      [36.53.Ø.195]
   All Campus                           IP      [36.255.255.255]
   Cerberus                             IP      [36.53.Ø.1Ø]
   Swanee                               IP      [36.56.Ø.2Ø8]
   Backbone A                           DLC     Ø207Ø1ØØ2BAF
   Backbone B                           DLC     Ø207Ø1ØØ2C6Ø
   Broadcast                            DLC     FFFFFFFFFFFF
   ClearView                            IP      [36.54.Ø.12]
   Fido                                 DLC     AAØØØ3Ø1131B
   Konig                                DLC     Ø26Ø8CØ3631Ø
   Tarpit                               IP      [36.8.Ø.47]
   Lundy                                IP      [36.54.Ø.11]
   pda                                  IP      [36.53.Ø.42]
            Use ↓ and ↑ then press ENTER, or ESC to return.
```

Figure 5–35. Dialog box to edit the working name table.

Each name consists of a pairing of *level* (that is, protocol) and *address*. For example, in Figure 5–35, the name *Fido* is attached to a station of level DLC and address AA0003 01131B. The name *Tarpit* is attached to a station of type IP and address [36.8.0.47].

3. To edit a name, move the highlight to the type-and-address pair you want to change, and press Enter.

   Result: The Sniffer analyzer opens a dialog box to receive a new name (Figure 5–36).

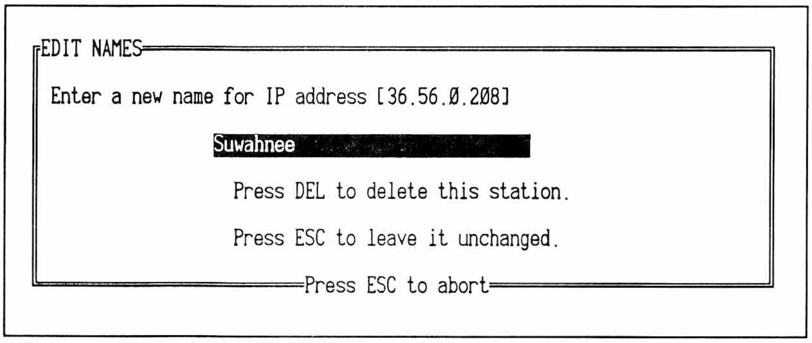

```
EDIT NAMES

Enter a new name for IP address [36.56.Ø.2Ø8]

                  Suwahnee

                  Press DEL to delete this station.

                  Press ESC to leave it unchanged.

                     Press ESC to abort
```

Figure 5–36. Dialog box to give a station a new name.

4. To add an address that is nowhere on the list, scroll to the top of the list to one of the lines marked **<new station>**. Each **<new station>** line is qualified by its level. There's always a line for the DLC level. In addition, there's a <new station> line for each level checked in the **address level** filter. Highlight the appropriate line and press Enter.

   Result: the Sniffer analyzer asks for an address and a name (Figure 5–37).[1]

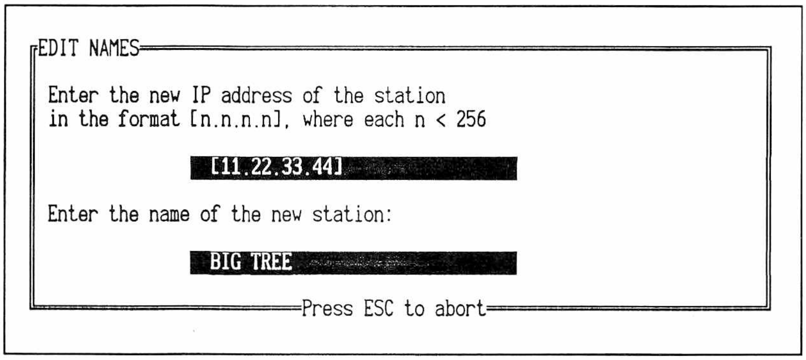

```
EDIT NAMES

Enter the new IP address of the station
in the format [n.n.n.n], where each n < 256

                [11.22.33.44]

Enter the name of the new station:

                BIG TREE

                     Press ESC to abort
```

Figure 5–37. Dialog box to type a new station's name and address.

[1] Although an IP address is always displayed with square brackets around it, during input, the analyzer accepts an address without brackets.

5. To prevent an address from being discarded, assign it a name. The name needn't be its final name. But an address whose name field is left blank is treated as an "unknown station" during capture, and is discarded from the name table when the table is rebuilt (at Save names, or when the analyzer is started).

When you specify a name for a station that already had a name, your new name thereby replaces the former name. Within the name table, addresses must be unique. However, names don't have to be unique; it's permissible to assign the same name to different addresses.

When you use the **Edit Names** dialog box, the Sniffer analyzer revises the name list *in working memory*. It doesn't update the stored file of names, so the revision is temporary. If you want to make the change permanent, you have to **Save names** (described next).

To save the current working name table so that it will be automatically effective next time you run the analyzer or monitor

In the main menu, select **Display**, then **Manage names**, and finally **Save names**. Press Enter.

Result: The analyzer copies all named addresses from the current working name table to the file c:\xxSNIFF\STARTUP.*xx*D.

Effect on the Sniffer Monitor's Name Table

The Sniffer analyzer and the Sniffer monitor make reference to the same name file. Like the analyzer, the monitor uses a working name table. When you start the monitor, it initializes its working name table with DLC addresses and their accompanying names copied from the saved names file STARTUP.*xx*D. The monitor ignores higher-level addresses. Thus, when you execute Save names, any changes in the names or DLC-level addresses will also affect subsequent uses of the monitor. Similarly, when the monitor saves names or addresses, those changes are merged into the file STARTUP.*xx*D and affect DLC-level names and addresses in the analyzer's subsequent displays.

Clearing the Name Table

The option **Clear all names** empties the Sniffer analyzer's working name table, removing both names and addresses. Use this command when you want to start over, perhaps before resolving names from a name file or looking for embedded names within higher-level protocols in the capture buffer.

Clear all names followed by **Save names** would empty not only the Sniffer analyzer's working name table but also the name table in the startup file.

Scanning the Capture Buffer for Addresses

The first time you display captured frames following a new capture, the analysis server scans the entire capture buffer for addresses. During the scan, it notes all addresses at any of the levels checked in the address level filter. When it finds a new address, it adds it to the working name table.

The name table has a maximum size. It has room for 500 addresses. During its scan for new addresses, the analysis server stops adding addresses when the table is full. Also, it never adds more than 50 addresses at any level, even when there's room for them

When the buffer contains many frames, the search may take a few seconds; the analysis server displays a thermometer-style gauge showing percentage completion.

Each address thus added has a blank name field. The analysis server sorts the address in alphabetical order by name. That puts the entries with blank names at the top of the list. When you edit names, it's easy to see where additional names are needed.

Taking Advantage of the Automatic Address Scan

When you expect to assign names to new high-level addresses, set the address level filter before you first display. That way the analysis server will automatically add up to 50 new addresses at each level; you have only to name them.

However, if you start display without setting the address level filters, you can still get the analysis server to add new addresses to the table. Press F6 for display options, go to address level, and check the levels you want. Press F3 to return to display. As you browse through the captured frames, whenever the Sniffer analyzer displays a frame, it adds any new addresses at the levels currently marked with a ✓. This doesn't scan the entire capture buffer, but it does add addresses from the frames you display.

Addresses you haven't named are temporary. When you execute **Save names**, the analyzer purges addresses that remain unnamed. So if you want to keep a record of new addresses, assign each a name — any name— before you execute **Save names.**

Looking for Names within the Captured Frames

Certain protocols (for example, Novell NetBIOS, or TCP DNS) let stations exchange tables of addresses and the names their users have adopted. When you ask the Sniffer analyzer to **Look for names,** it scans the capture buffer for such messages. From them, it copies two kinds of entries to the working name table:

- Names for addresses already in the name table without names;
- Both name and address for addresses not yet in the table.

This technique may find names or addresses for stations that weren't themselves sending or receiving.

To look for names within the frames in the capture buffer

1. Capture some frames, live from the network or by replay from a file.
2. Press F3, **Display**.
3. Press F6, **Display options**. Move the highlight to **Manage names** and then to **Look for names**. Press Enter.

   Result: The analyzer scans the entire capture buffer for protocols that exchange name information, and adds new addresses and their names to the working name table. (It does not revise names for addresses that were already in the table.)

   The analyzer reports the number of address-and-name pairs it added to the working name table.

Names found this way may be quite transitory. For example, you may find a name that a user assigned for a single work session. The user may move to another machine and assign the name to a different address. Because such names are so readily changed, you may not want to save a name table constructed this way; it may be wrong next time you use it.

Resolve Names by Referring to Other Name Files

To help identify unnamed addresses, you may have the Sniffer analyzer look up names from an external file. The file from which names are taken is in the same format as STARTUP.*xx*D. Presumably it exists because at some earlier time you copied and renamed the version of STARTUP.*xx*D that was then current, or downloaded it to the server from the console.

To resolve unnamed stations by searching an external file

1. Capture some frames, live from the network or by replay from a file. (The procedure won't run if there's nothing in the capture buffer.)
2. Move the highlight to **Display filters**, then **Address level.** Put a check mark at each level you want included in the search.
3. Move the highlight to **Display**, then **Manage Names,** then **Resolve Names**. Press Enter. (It doesn't matter whether you first display the capture buffer.)

4. Press F6, **Display options**. Move the highlight to **Manage names** and then to **Look for names**. Press Enter.

   Result: The Sniffer analyzer opens a dialog box showing a list of files that it recognizes as name files. Move the highlight to the one you want and press Enter.

   The analyzer scans the entire capture buffer for addresses at the DLC level and at any higher levels you checked in the **Address level** filter. The analyzer amends the working name table by adding to it the address and level of any station represented in the capture buffer but not named in the name file. That gives it a pooled list of unnamed stations, including unnamed stations that it just found in the capture buffer and any that were in the name table but unnamed.

   For each address in its list of unnamed stations, it searches the external file. Wherever it finds a name for one of its unnamed addresses, it inserts the name in the name table.

   When the analyzer has completed its search, it displays a message telling you how many unnamed addresses were in the working name table, and how many of them it was able to name.

5. To preserve the names you have added, execute **Save names** to copy the working name table into the startup file. This saves addresses that have names, but discards addresses that still lack names.

Importing a Name File

You can upload files from the SniffMaster console to the server. That gives you a way to install name files that you acquired from elsewhere, or that you constructed with a text editor. Once you have the additional name files, the **Resolve names** command can search them for names to go with the as-yet-unnamed addresses in the server's working name table.

For the analyzer to recognize and use a name file, the file's name must have an identifying extension and its text must be in a standard format. The name file's name and format are described in Chapter 6, see pages 6–10 and 6–14.

For the procedure to upload a file from console to server, consult *Distributed Sniffer System: Installation and Operations Manual.*

Total Size of the Name Table

The Sniffer analyzer sets a limit on the size of the name table. The overall limit is 500 entries.

There's an additional limit of 50 names at each address level each time you start a new display or each time you run **Resolve names**.

Files of Saved Frames

The Sniffer analyzer permits you to save all or part of the contents of the capture buffer to a file. Or, you can load the capture buffer with frames from a previously-saved file. By loading previously-saved frames, you can examine frames captured at another time or place, or frames sent to you for study.

Loading a File of Previously-Saved Frames

You can load the capture buffer directly from a previously saved file, without going through capture.

To load the capture buffer with a file of saved frames

1. In the main menu, move the highlight to **Files**, then **Load**, and then **Data**, and press Enter.

   Result: The Sniffer analyzer opens a dialog box containing a list of previously-saved capture files and directories. The list is in alphabetical order. A capture file has an extension that starts with a two-letter code for the network (EN for Ethernet, TR for token ring, SY for synchronous WAN). The third letter is C (for capture). The analyzer only shows you filenames whose extensions indicate that they're from the network for which your server is configured.

2. To change to a different directory, move the highlight to one of the rows labeled **<DIR>**.

   The notation **.. <DIR>** in the top row indicates the directory one step nearer the root. Other entries with **<DIR>** are subdirectories. To see the list of files in a subdirectory, highlight its name and press Enter.

3. Move the highlight to the file you want. Alternatively, to jump directly to a part of the list, type a letter from the keyboard. The highlight moves to the next entry starting with that letter.

   When you've highlighted the name of the file you want, press Enter to load it.

4. Press F3 to display the capture buffer.

Saving a File of Frames

While you have frames in the capture buffer (either because you just captured them, or because you loaded them from a file), you can save the frames in the buffer to a file. You can save:

- Everything in the capture buffer, or
- Only those frames that pass your current display filter; or
- Only frames within a particular range.

To save frames in the capture buffer to a file

1. While displaying the capture buffer, if you wish to save only those frames that pass the display filters, set the filters appropriately (see "Setting the Display Filters" on page 5–5ff).
2. Press F5 to return to the main menu. Move the highlight to **Files**, then **Save**, then **Data**. Don't press Enter yet.

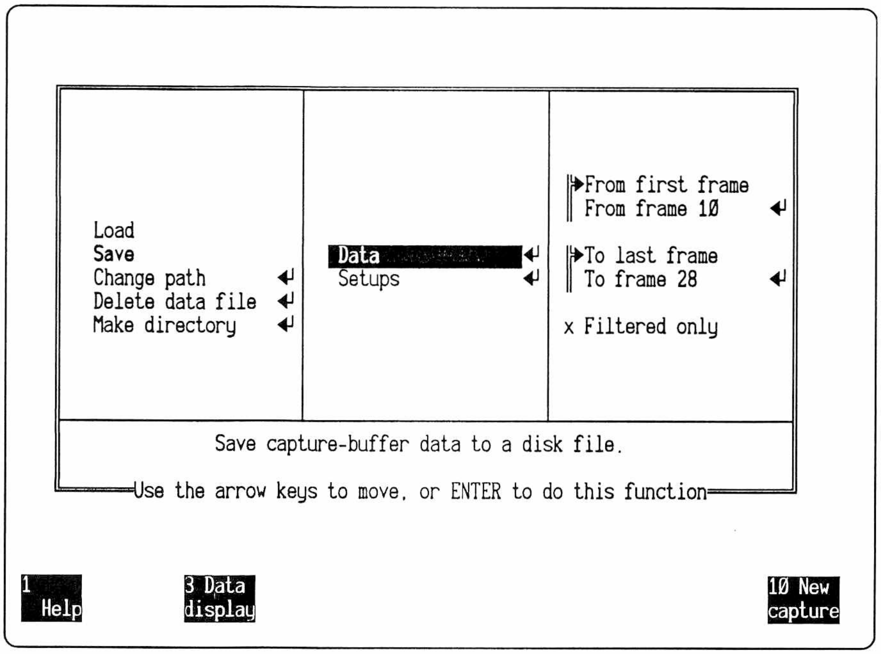

Save capture-buffer data to a disk file.

Use the arrow keys to move, or ENTER to do this function

Figure 5–38. Menu to save captured frames to a file.

3. Set the range of frames to be saved:
   - First frame
   - Last frame

   To change the frame number, move the highlight to it and press Enter. The analyzer opens a dialog box to receive a new number. Type the number you want and press Enter. (It's the same procedure as setting the range of frames to be printed; see "To send a report to file or printer" on page 5–57).

4. Indicate whether you want to save all frames or just those that pass the current filters. Highlight **Filtered only** and press Spacebar to toggle between ✓ (active) and X (inactive).

5. When you've completed the specification (except the filename) move the highlight back to **Data** and press Enter.

   Result: The analyzer opens a dialog box for you to write a name for the file to be created. Initially, the dialog box contains the path but no name. (To change the default path, or to create a new directory, see page 6–9.)

   The path shows C (the server's hard drive) and the current directory. You can backspace over that part of the display and write the name of a different directory if you wish.

   Write your filename using no more than eight characters. Don't include an extension; the Sniffer analyzer automatically attaches the extension .xxC (where xx is the two-letter network code, EN for Ethernet, TR for token ring, or SY for WAN/ synchronous). If the name you propose is already in use, the analyzer warns you; you can abort the request or go ahead and overwrite the existing file.

Deleting a File of Saved Frames

The **Files** menu also offers you the option to delete a file of saved frames. When you highlight **Delete** and press Enter, the analyzer opens a dialog box showing the names of files of saved frames, in much the same way as it shows a list of files you could load or replay (see "To capture from a file" on page 2–50). The list includes only capture files related to the network you are now using.

Pressing Enter while the name of a file is highlighted indicates that you want to delete it. The Sniffer analyzer opens a dialog box headed **Warning** with the message

`Deleting:` *xxxxxxx*

where *xxxxxxx* is the name of the file.

You can press Enter to go ahead and delete the file, or Esc to return to the list of data files without deleting it.

Saved Setups

Your "setup" is the particular constellation of display options, filters and controls you're using. You can save the setup to a file. Later, you can load that file, thereby restoring all those options to the values they had when you saved the setup.

To save your current setup

1. Check that the way things are now is the way you want them to be restored. (See the next section for a listing of what's included in a setup and what's not.)

2. In the main menu, move the highlight to **Files,** then **Save,** then **Setup**. Press Enter.

   Result: The analyzer opens a dialog box for you to write the file's name. Initially, the dialog box contains the path but no name. (To change the default path, or to create a new directory, see page 6–9.)

   The path shows C (the server's hard drive) and the current path. You can backspace over that part of the display and write the name of a different directory if you wish. Write your file name using no more than eight characters. Don't include an extension; the Sniffer analyzer automatically attaches the extension .*xx*S (where *xx* is the two-letter network abbreviation). If the name you propose is already in use, the analyzer warns you; you can abort the request or go ahead and overwrite the existing file.

Contents of a Setup File

The setup file records every option that you can set by a check mark or a radio control. Some items set in other ways are also included, but some are not.

Included in a setup file

- General options (for example, WAN/synchronous encoding, token ring remove-if-no-signal).
- Capture options (for example, whether to display skylines or meters, whether to show individual or pairwise counts).
- Capture filters (including source and destination addresses, and pattern match logic together with the individual patterns and their offsets).
- Trigger options (including options for stopping capture and pattern match logic together with the individual patterns and their offsets).
- Display options (including views and levels).
- Display filters (including source and destination addresses, and pattern match logic together with the individual patterns and their offsets).

- Printer options (including whether to send output to a printer or to a file, and whether to include headings and page numbers).

***Not* included in a setup file**

- The path you may have set for locating the directory of files to be loaded or saved. However, you can record that with the **Set path** command.
- The capture buffer. However, you can save that by the command sequence **Save** and then **Data**.
- The name table. However, you can save the current working name table by the command sequence **Manage names** and then **Save names**.

  Your setup may refer to addresses (for example, in a station address filter). When you display the filter, you may see the station's symbolic name rather than its numeric address. However, the saved filter really contains the numeric address. The symbolic name is regenerated during display by looking for that address in whatever version of the name table is then current.

Using a Saved Setup File

When you first start the Sniffer analyzer, it checks for the existence of a file called STARTUP.xxS. If it finds such a file, it treats it as a setup file and loads it automatically. Otherwise, it uses the Sniffer analyzer's standard defaults.

To activate a saved setup manually

1. In the main menu, move the highlight to **Files**, then **Load**, then **Setup**, and press Enter.

   Result: The Sniffer analyzer opens a dialog box showing you a list of saved setup files. Move the highlight to the one you want and press Enter.

   The analyzer loads the setup file. There is no specific message saying it has done so. It resets all options the way the setup file specifies.

To restore Network General's default options

1. Follow the procedure for loading a saved option file. When the analyzer prompts you for the file's name, select the file called c:\CAPTURE\DEFAULT.*xx*S.

CHAPTER SIX: THE ANALYSIS SERVER'S USE OF FILES 6

Chapter 6. The Analysis Server's Use of Files

Chapter Overview

This chapter describes:

- Files that the Sniffer analysis server uses and the directories that contain them.
- The internal format of files used to store name tables, setups, and saved traces.

During everyday use of the Sniffer analysis server, you should have little or no need for the information in this chapter. However, knowledge of the file organization may prove useful when you import files from somewhere else or export files produced by the analyzer for analysis elsewhere. This chapter presumes that you're familiar with the naming and storage conventions of the DOS file system.

File Names

Under the DOS operating system, a file's name is composed of two parts. The two parts are separated by a dot. The first part (sometimes called the base name) may contain up to eight characters. The second part (called the extension) consists of three letters. (Omitting the dot and the letters that follow it has the same effect as writing a dot and three blanks.)Chapter 6, "Chapter Figure 6–1The Analysis Server's Use of Files."

Certain extensions have special significance to DOS. For example, an executable file has the extension .EXE, .COM or .BAT; DOS won't execute a file whose name has some other extension. Figure 6–1 shows some of the extensions that occur in files used by the Sniffer analysis server. The first two of these are standard DOS conventions, while the others are more specific.

Extension	File type
.EXE	Executable file. The extension may be omitted from the DOS command that invokes its execution.
.BAT	Batch file (an executable file consisting of commands in the DOS shell language). The extension may be omitted from the DOS command that invokes its execution.
.PRN	Output file generated by a report generator for printing (whether or not sent directly to a printer).
.CSV	Output as "comma-separated values," intended as input to a spread sheet.
.TXT	Script used to generate the selection menu; used by the executable file MENUX.EXE.
.MNU	Menu script to generate a particular analysis entry in the selection menu.
.CFG	Configuration script describing the protocol interpreter suites for a network.
.HLP	Help file to explain choices in the Sniffer analyzer's menus.

Figure 6–1. Extensions in Sniffer file names.

Each of the analyzer's executable files refers to files that are specific to a single network. Each of these files has a name whose extension classifies it by use and by network. The first two letters of the extension indicate the network, and the last letter indicates the type of file. For example, if you're working on Ethernet and save the capture buffer to a file called MYDATA, the analyzer gives it the extension .ENC (Ethernet captured frames). Letters used in extensions are listed in Figure 6–2.

First Two Letters	
EN	Ethernet
TR	Token Ring
SY	WAN/Synchronous.

Last Letter	
C	Captured frames.
S	Setup (values of options used in operating the analyzer).
D	Station names (file of symbolic equivalents for numeric addresses).
I	IDs (symbolic equivalents for the manufacturer codes within numeric addresses).
T	Type (used by Sniffer Network Monitor for symbolic equivalents for types of DLC frames).
B	Binary type (used by the Sniffer monitor to summarize the setup of the user interface).

Figure 6–2. Letters used in the extension of an analyzer file's name.

Names for the Sniffer Analyzer's Executable Files

You don't ordinarily see the name of the executable files that in fact provide the Sniffer analyzer's services. That's because, to start work, you move the highlight to an entry in the selection menu; you don't invoke the file by typing its name. An item in the menu describes a set of services (for example "Token Ring Analyzer") but doesn't show you the name by which DOS identifies the underlying executable file. For each service listed in the menu, there's a separate executable file. Each executable file for a Sniffer analyzer has a name that encodes the network and the protocol interpreter suites it provides. The name is constructed as follows:

The first two letters indicate the network, as shown in Figure Figure 6–2. Those two letters are always followed by the letters SN.

The next four positions are used to encode the protocol interpreter suites that are linked into the executable file. This encoding uses any of the 32 characters 0–9 or A–V. This base-32 encoding permits any combination of protocol suites to be given a unique name. For example, analysis for Ethernet with interpreter suites for DECnet, TCP/IP and ISO has the names ENSN2O00.EXE. Analysis for token ring with the same suites and also the IBM protocol suite has the name TRSNIO00.EXE.

Directories

Figure 6–3 summarizes the principal directories on the Sniffer analysis server's hard disk (drive C).

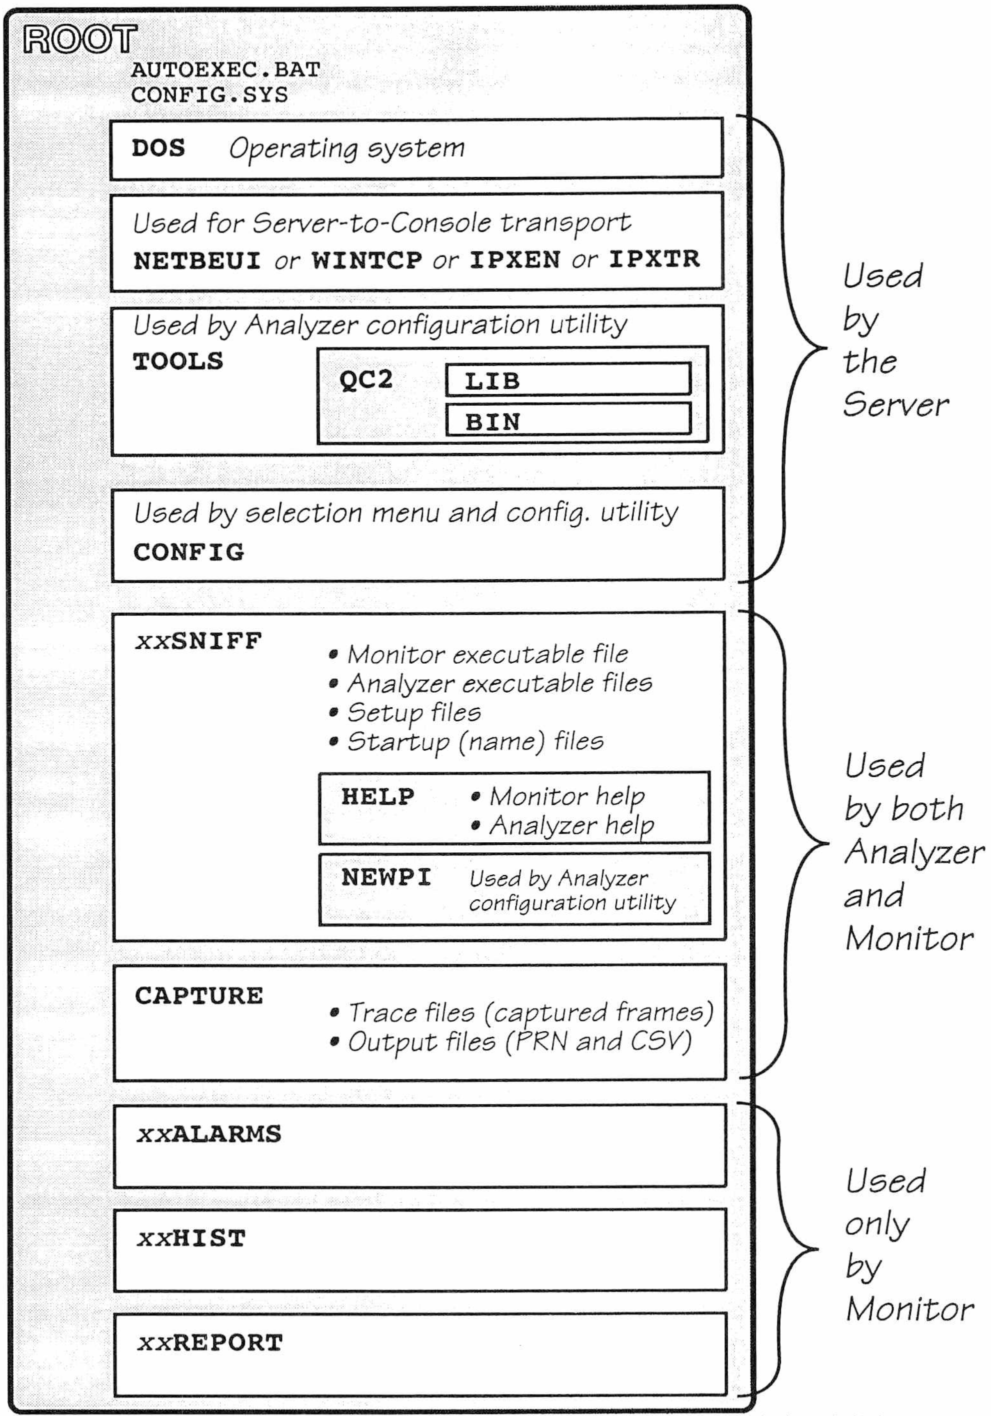
ROOT
AUTOEXEC.BAT
CONFIG.SYS
DOS Operating system
Used for Server-to-Console transport
NETBEUI or WINTCP or IPXEN or IPXTR
Used by Analyzer configuration utility
TOOLS
QC2
LIB
BIN
Used by selection menu and config. utility
CONFIG
Used by the Server
xxSNIFF
• Monitor executable file
• Analyzer executable files
• Setup files
• Startup (name) files
HELP
• Monitor help
• Analyzer help
NEWPI
Used by Analyzer configuration utility
CAPTURE
• Trace files (captured frames)
• Output files (PRN and CSV)
Used by both Analyzer and Monitor
xxALARMS
xxHIST
xxREPORT
Used only by Monitor

Figure 6–3. Directories on the Sniffer analysis server's hard disk.

The root directory contains the files AUTOEXEC.BAT and CONFIG.SYS. The server refers to them automatically each time you reboot the server or restart it following a power-down.

The Sniffer analysis server's AUTOEXEC file establishes the path: that is, the list of directories in which the operating system searches for executable files. The AUTOEXEC file invokes the selection menu. From there, you select one of the Sniffer analysis server's major functions (Monitor, Analyzer, File Transfer, and so on).

DOS Directory

The DOS directory contains files belonging to the operating system (except that a few DOS files that are in the root directory). The path includes the DOS directory, so the operating system looks there to find its own files. Normally, you won't need to make any changes to files in the DOS directory.

CONFIG Directory

The CONFIG directory contains scripts that generate the Sniffer analysis server's selection menu, the analyzer's main menu, the monitor's main menu, and the menu used by the configuration utility.

Menu files

Each executable file is accompanied by a menu file whose name is the same but has the extension .MNU rather than .EXE. The executable files reside in the xxSNIFF directory (that is, ENSNIFF, TRSNIFF, or SYSNIFF). However, all .MNU files reside in the \CONFIG directory. For example, the executable file TRSNIO00.EXE (in the directory \TRSNIFF) is accompanied by the menu file TRSNIO00.MNU (in the \CONFIG directory).

The server uses the various menu files to generate entries in the main selection menu. Each menu file is responsible for the menu entry corresponding to one analyzer or monitor. The .EXE files and the .MNU files are already supplied by Network General. When the configuration utility (described in the next chapter) generates a new executable file, it automatically generates the corresponding .MNU file as well. Similarly, when the utility deletes an executable file, it also deletes the corresponding .MNU file.

Configuration Files

The configuration files list the facilities available to Sniffer analyzers on the particular server. Its name is formed from the two-letter network abbreviation, the letters SNIFF, and the extension .CFG (for example, for Ethernet, ENSNIFF.CFG). The .CFG file lists all the protocol interpreter suites that can be installed in an Ethernet analyzer. When you use the Configuration Utility to build an analyzer with a new combination of protocol interpreter suites, the menu you see (listing the available interpreter suites) is governed by information in the configuration files.

Tools Directory

The TOOLS directory contains the program that generates the Sniffer selection menu, and various other utilities that go with it. The Sniffer server's DOS path directs the operating system to the \TOOLS directory, so you won't usually need to make any explicit reference to it.

The \TOOLS directory has a subdirectory QC2. It contains the Microsoft Quick-C compiler, and its subdirectories LIB and BIN. The configuration utility makes use of these, but you never have to deal with them explicitly.

*xx*SNIFF Directory

One directory contains the principal executable files, both for the monitor and for the various analyzer configurations. The directory's name is formed from the two-letter network abbreviation followed by the letters SNIFF. For example, for token ring, the directory is called TRSNIFF; for Ethernet it's ENSNIFF; and for WAN/Synchronous, it's SYSNIFF. A server has whichever of these directories is appropriate, but not more than one of them.[1]

Capture Directory

The \CAPTURE directory is where the Sniffer analyzer ordinarily saves files of captured frames. It's also the default destination for output files such as .PRN files of .CSV files.

Creating Alternate Directories for Saved Files

To keep various sets of captured data separate (for example, data collected at different times or under different circumstances), you may want to set up additional directories to contain the files. The batch file that starts the Sniffer analysis server makes \CAPTURE the current directory.

[1] A self contained portable Sniffer analyzer may be equipped for more than one network, in which case it has more than one such directory.

To create a new directory

1. In the main menu, move the highlight to **Files**, then to **Make directory**, and press Enter.

   Result: The analyzer opens a dialog box to receive the name of the new directory. The analyzer supplies the current path. If the path hasn't been changed, it is C:\CAPTURE\.

2. Write your new name to the right of the final \. That will make your new directory a subdirectory of \CAPTURE. Alternatively, you can backspace into all or part of the name the analyzer suggested, and replace it with whatever you prefer. To record the name and create the directory, press Enter.

The analyzer limits the name to alphabetic characters, so you can't specify an extension for a directory created in this way. The operating system does not accept more than 8 characters in the main part of a name. If you write a longer name, the analyzer truncates it to the first 8 characters. The Sniffer analyzer accepts the path you specify. It doesn't verify that your path is syntactically valid or that the directory actually exists. If the path is not valid, when you subsequently try to read or write a file, you will get the error message `Invalid path.`

Setting the Path to a Directory for Saved Files

Whenever the analyzer opens a dialog box in which you select an existing file for input or name a file for output, it starts by showing you the path to its current directory. Initially, the path is

C:\CAPTURE\

To specify the initial path to saved files

1. From the main menu, move the highlight to **Files**, then to **Change path**, and press Enter.

   Result: The analyzer opens a dialog box in which you can edit the path in accordance with your wishes.

2. Write the path you want. The path should end in \. To record your preference, press Enter.

   Result: Whenever it opens a dialog box for you to specify a file to read or write, the analyzer will start from the path you recorded here.

The analyzer accepts the path you specify. It doesn't verify that your path is syntactically valid or that the directory actually exists. If the path is not valid, when you subsequently try to read or write a file, you will get the error message `Invalid path.`

Names for Files of Captured Frames

A data file created by saving frames from the capture buffer is called a *capture* file (or sometimes a *trace* file). It has the same internal format as the capture buffer. You cannot read a capture file as text.

When you create a data file, assign it any name you wish in eight letters or fewer. The Sniffer analyzer automatically supplies the name's extension (the three letters following the dot). The extension consists of the two-letter network code followed by the letter C for "capture." The codes are shown in Figure 6–4.

Ethernet	.ENC
Token Ring	.TRC
WAN/ Synchronous	.SYC

Figure 6–4. Extensions in the names of capture files.

Storage of Setups and Name Tables

For information about names or setups, the Sniffer analyzer refers to three kinds of files. Each type has a characteristic extension. As described in Figure 6–2 (page 6–5), the extension consists of the two-letter network abbreviation (EN for Ethernet, SY for synchronous, TR for token ring) and a final letter to indicate the type of file. Each time the Sniffer analyzer starts execution, it refers to one or more of these. The types of files are listed in Figure 6–5.

File Name	Contains	How Used
STARTUP.xxD	Symbolic names for addresses.	Each time you start the analyzer, it builds its working name table by reading names from \xxSNIFF\STARTUP.*xx*D. You may edit the working name table (**Manage names** then **Edit names**) and save the revision (**Manage name** then **Save names**).
STARTUP.xxI	Symbolic names for manufacturer ID codes	Each time you start the analyzer, it reads the table of manufacturer IDs from \xxSNIFF\STARTUP.*xx*I.
STARTUP.xxS *or* *(after startup)* ANYNAME.xxS	Values of user-settable options in the analyzer menus.	None is required and none is initially provided. *If* you have saved a setup file and given it the name \xxSNIFF\STARTUP.*xx*S, the analyzer automatically loads it at startup. (Other setup files are usually saved in and loaded from the \CAPTURE directory.)

Figure 6–5. Files referred to at startup.

Each time you start a Sniffer analyzer, the software automatically checks for the presence of the startup files it needs. It looks for them in the \xxSNIFF directory. If it finds them, it uses them to set its working name table, or to initialize the analyzer's filters and settings. If the analyzer doesn't find a file it needs, it uses a default setup or an empty name table.

Saving Setups and Names to a File

The three types of files (name files, manufacturer ID files, and setup files) have different mechanisms for loading and saving. These mechanisms are summarized in the paragraphs that follow and in Figure 6–6. In the figure, a gray arrow indicates automatic loading each time the analyzer is started. A black arrow indicates a file that you can load or save upon command.

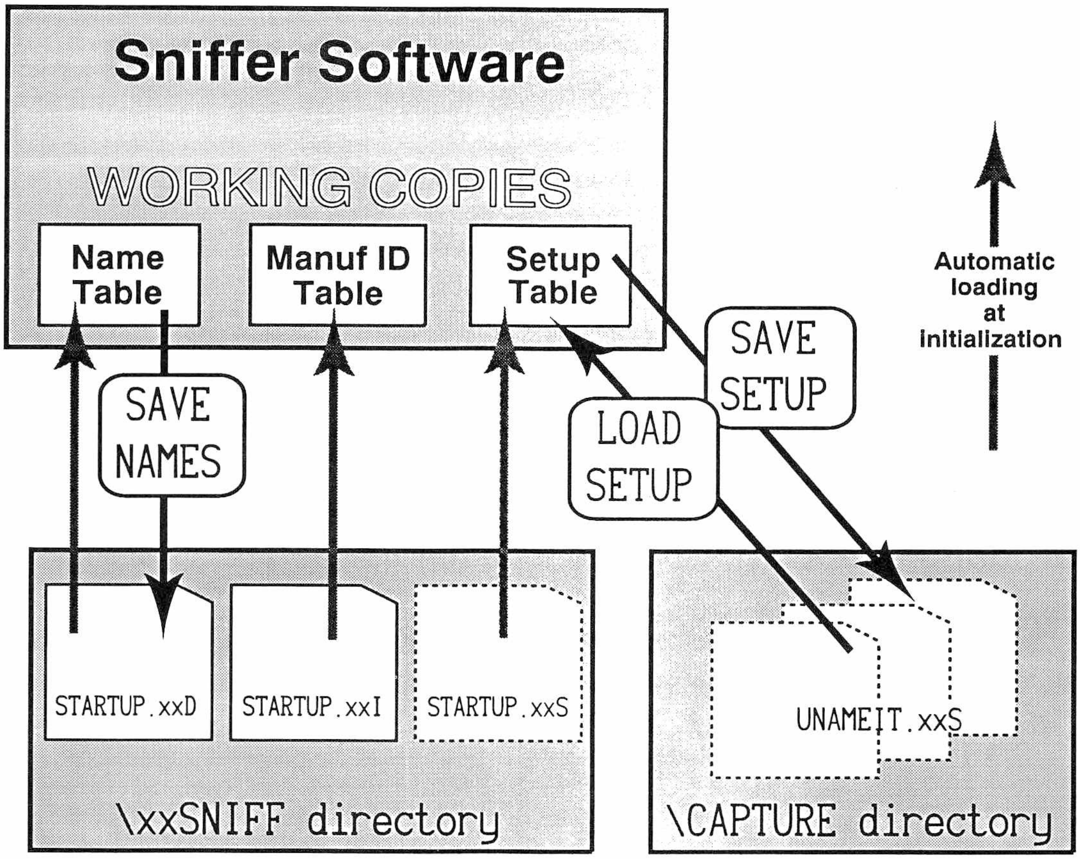

Figure 6–6. Stored files and working copies of name tables and setups.

Modifying the Default Setup

The Sniffer analyzer does not require a file called STARTUP.*xx*S. When the analyzer starts, it checks to see whether such a file exists. If it exists, the analyzer uses it. If it doesn't exist, the analyzer uses its own list of default settings for options.

To alter the setup with which the analyzer starts

This is a variant of "To save your current setup" on page 5–71.

1. A saved setup is a snapshot of the analyzer's current settings. You should start by making sure that your current settings are the ones you would like to have the analyzer supply automatically at startup. (For a listing of what's included in a setup and what's not, see "Contents of a Setup File" on page 5–71.)
2. In the main menu, move the highlight to **Files**, then **Save**, then **Setup**. Press Enter.

   Result: The analyzer opens a dialog box for you to write the file's name. Initially, the dialog box contains the path but no name.

   Overwrite whatever is now in the dialog box. Replace it by
   C:\ENSNIFF\STARTUP (for Ethernet)
   C:\TRSNIFF\STARTUP (for token ring)
   C:\SYSNIFF\STARTUP (for WAN/synchronous)
   Don't include an extension; the Sniffer analyzer automatically attaches the extension .*xx*S (where *xx* is the two-letter network abbreviation). If the file already exists, the analyzer warns you; you can abort the request or go ahead and overwrite the existing file.

Restoring the Setup to Network General's Defaults

Suppose you have been using a customized setup and wish to revert to the default setup with which Network General shipped the Sniffer analysis server. You can do that in either of two ways, one lasting and the other temporary.

To make a lasting return to Network General's default setup

1. From the analyzer's main menu, highlight **Exit** and press Enter
2. In the server's selection menu, highlight **Exit to the operating system** and press Enter.
3. At the DOS prompt, type the command to rename the file now called *xx*SNIFF\STARTUP.*xx*S.

Choose a new name that preserves the extension *.xx*S but changes STARTUP to anything else. The DOS command is

```
RENAME STARTUP.xxS ANYTHING.*
```

Result: Each time you start the analyzer thereafter, it will find no file called STARTUP.*xx* S, and so it will use its standard defaults.

4. To return to the Sniffer server's selection menu, type

```
menu
```

and press Enter.

To make a temporary return to Network General's default setup

1. From the analyzer's main menu, move the highlight to **Files**, then **Load**, then **Setup** and press Enter.

   Result: The analyzer opens a dialog box in which you may choose the name of a setup file to load.

2. Choose the setup file called \CAPTURE\DEFAULTS.*xx*S and press Enter.

   Result: The file contains a duplicate of the analyzer's default settings. When you load this file, it restores the default setup. This will *not* affect the initial setup the next time you start the analyzer.

Name Tables

Procedures for assigning names to stations are described in Chapter 5, starting at page page 5–62. This section describes the format of files that contain name tables.

While it runs, the Sniffer analyzer uses an internal directory called the *working name table*. Each time you start the analyzer thereafter, it initializes the working name table by reading from a file. Ordinarily, the file it reads is *xx*SNIFF\STARTUP.*xx*D.

If the analyzer can't find *xx*SNIFF\STARTUP.*xx*D, it looks in the current directory (that is, the directory identified by the Change path command). The batch file that starts the server normally makes \CAPTURE the default directory, so the analyzer looks next for \CAPTURE\STARTUP.*xx*D.

If it doesn't find that file either, the analyzer sets up an empty name table, containing only the address of the server's own monitoring card and the name "This Sniffer."

You may have additional reference files of names and station addresses. These are either renamed copies of what was once a STARTUP.*xx*D file, or files in the same format created by an editor and downloaded to the server. When you create such a file, the name you give it must be something other than STARTUP, but must have the same extension as a startup file.

You can have the analyzer consult these additional name files by executing the **Resolve names** command (page 5–66).

In all cases, the extra name files are used as a source of names to fill blanks in the working name table. There is no command to load an entire substitute name file the way you load a setup.[1]

Building Name Files

A name file is identified by an extension consisting of the two-letter network code followed by the letter D. There are two principal ways to generate a name file: by saving and then renaming the working name table, and by writing a name file from scratch with a text editor and then transferring it to the server.

To build a name file from the current working name table

1. From the analyzer's main menu, move the highlight to **Display**, then **Manage names**, then **Edit names**, and press Enter.

   <u>Result</u>: The analyzer opens a scrollable dialog box in which you can edit the name table.

2. Put the information you want into the working name table (see "To edit the working name table" on page 5–62). When you've finished adjusting the table, press Esc to return to the menus.

3. Move the highlight to **Save names** and press Enter. The analyzer saves a file called \xxSNIFF\STARTUP.xxD. It contains names and address for all stations that had been named in the working name table, but omits stations to which you had not assigned names.

4. Exit from the Sniffer analyzer. At the server's selection menu, highlight **Return to the operating system** and press Enter.

5. At the DOS prompt, copy and rename the file STARTUP.xxD. Put the new copy into a convenient place (for example, the \CAPTURE directory). Give it a name that is not STARTUP but has the same extension, by a command such as

   ```
   COPY \xxSNIFF\STARTUP.xxD \CAPTURE\newname.*
   ```

[1] If you wish to maintain several independent name files, use DOS commands to give them arbitrary names. Before you start the analyzer, (a) assign a new name to your existing file STARTUP.xxD, and then (b) copy the one you wish to make active, giving the copy the name STARTUP.xxD.

Writing Your Own Name File

Alternatively, you can build your own name files directly. This section describes a name file's internal format.

All name files have the same format, which applies both to the file called STARTUP.*xx*D and any other name files you may use as sources. Figure 6–7 shows part of a name file. Since a name file is a standard ASCII file, you can built it at the console or any other PC using any standard text editor.

```
station "Broadcast"    = addrtype "DLC" C000FFFFFFFF
station "Error Log"    = addrtype "DLC" C00000000008
station "ipS1"         = addrtype "IP" [36.10.0.13]
station "ipS2"         = addrtype "IP" [36.11.0.14]
station "ipS3"         = addrtype "IP" [36.11.0.23]
station "ipS4"         = addrtype "IP" [36.2.0.5]
station "ipS5"         = addrtype "IP" [36.22.0.20]
station "Long, 31-Character Station Name" = addrtype "DLC" 400000000002
station "Mary"         = addrtype "DLC" 10005A0033BF
station "This Sniffer" = addrtype "DLC" 48000A000001
station "Tom"          = addrtype "DLC" 10005A002FEB
station "Faquard"      = addrtype "XNS" 08000AC7CEFE
```

Figure 6–7. Sample name file.

For convenience during subsequent display, the **Save names** command sorts the rows of the table alphabetically by name. However, the analyzer does not require a name file to be in alphabetical order.

The following features of a name table are illustrated in Figure 6–7 and Figure 6–8.

Type

Each line starts with a word or symbol that identifies its type. The three types are distinguished by the following as their first non-blank characters:

`station` The balance of the line describes one station's name and address.

`addrtype` The balance of the line sets the default address type (protocol) that applies to subsequent lines that don't include it explicitly.

/* A line that starts with /* and ends with */ is a comment and is not executed.

Name

The station's name to the right of the word station. The name must be enclosed in double quotes.

Address type

Each address must be assigned to a specific type (that is, protocol). The type can be stated in either of two ways:

- Explicitly for each address, by including the phrase
  `addrtype "DLC"`
  to the left of the address (as shown in Figure 6–7). A name file generated by the **Save names** command is entirely in this form.
- Implicitly, using the current default type (Figure 6–8).An address that has no explicit type is presumed to belong to the current default type. The default type is initially "DLC". The default type is set by each use of **addrtype**, and remains in effect until another **addrtype**.

Address

There is an = sign between the name and the address. An address is not enclosed in quotes. Each address is written in a form appropriate to its type (the same way the Sniffer analyzer displays it in the detail view). For example, a 6-byte DLC address is written as 12 hexadecimal digits. A 4-byte IP address is written as four decimal numbers separated by dots and entirely surrounded by square brackets.

Name Table with Default Types

Figure 6–8 shows the same information as Figure 6–7, but makes use of *default types*. To improve readability, the file may contain redundant blanks or blank lines, as well as comments. A comment starts with /* and continues to */.

```
 station "Error Log" = CØØØØØØØØØØ8
addrtype "IP"
 station "ipS1" = [36.1Ø.Ø.13]
 station "ipS2" = [36.11.Ø.14]
 station "ipS3" = [36.11.Ø.23]
 station "ipS4" = [36.2.Ø.5]
 station "ipS5" = [36.22.Ø.2Ø]
 station "ipS6" = [36.26.Ø.54]
 station "ipS7" = [36.26.Ø.56]
addrtype "DLC"
/* Long name inserted as a test */
 station "Long, 31-Character Station Name"=4ØØØØØØØØØØ2
 station "Mary" = 1ØØØ5AØØ33BF
 station "This Sniffer" = 48ØØØAØØØØØ1
 station "Tom" = 1ØØØ5AØØ2FEB
addrtype "XNS"
 station "Faquard" = Ø8ØØØAC7CEFE
```

Figure 6–8. The name file of Figure 6–7 rewritten with default types.

Alphabetization of Station Names

The order of entries in a name file doesn't matter.

When the Sniffer analyzer builds and displays its working name table, it shows the list in alphabetical order by name. Addresses that you haven't named (and therefore have blank names) appear at the top of the list.

Each time you edit the working name table, the Sniffer analyzer re-alphabetizes the list. When you execute **Save names**, the saved file preserves the order of the working name table (but discards entries that aren't named).

Table of Manufacturer ID Codes and Abbreviations

On most LANs, an address consists of six bytes. The first three represent the manufacturer. The analyzer attempts to represent the first three bytes by a six-character abbreviation of the manufacturer's name. The address then appears as six characters of manufacturer abbreviation followed by six characters of hexadecimal, for example Intrln031EF7.

The table of manufacturer codes and names is located in the file startup.*xx*I, where *xx* is the two-letter code for the network, and I indicates the ID table. The file's internal format is illustrated in Figure 6–9. The figure shows the content of the file STARTUP.*xx*I for a network that transmits least-significant-bit first (for example, Ethernet but not token ring). The comments in the file are for the convenience of human readers and have no effect on the analyzer's use of the file. (For compactness, the lower part of the table is shown here with three entries per line. In the file, each entry has a line of its own.)

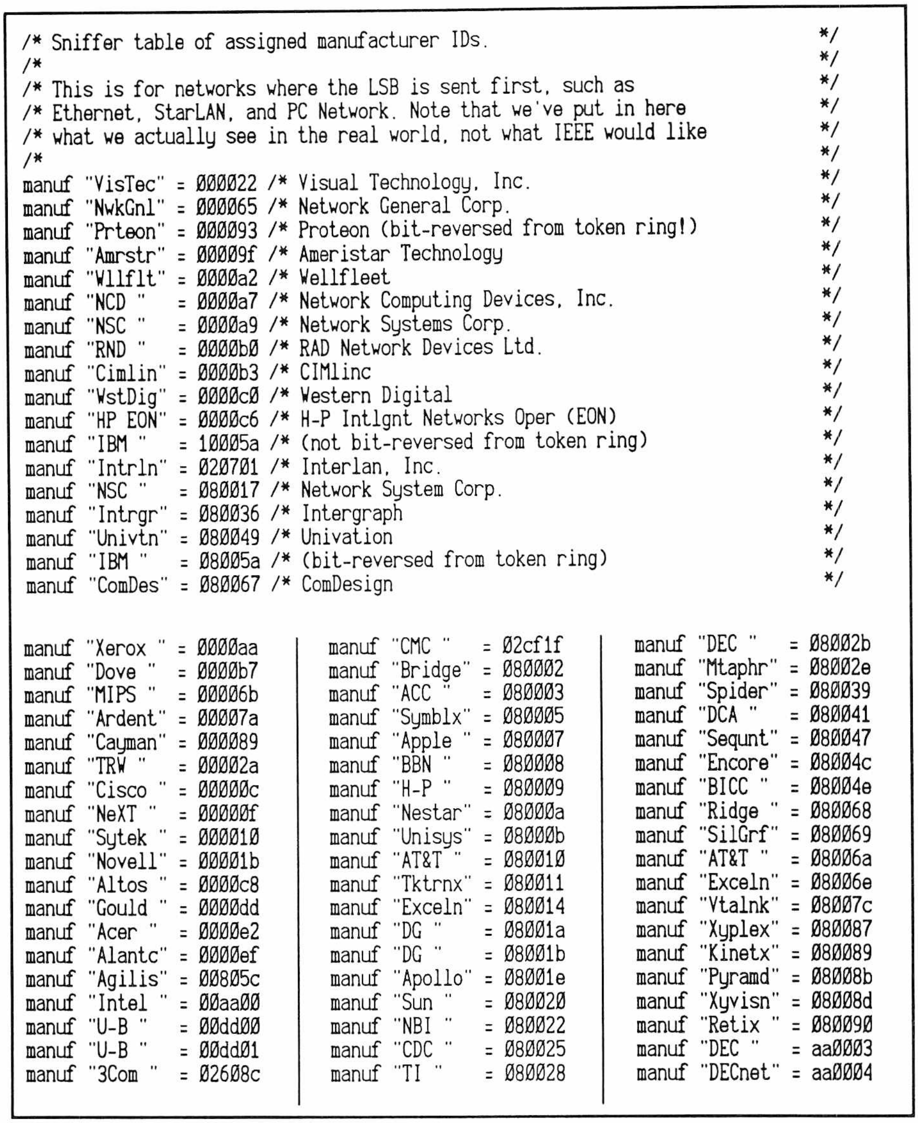

```
/* Sniffer table of assigned manufacturer IDs.                                   */
/*                                                                               */
/* This is for networks where the LSB is sent first, such as                     */
/* Ethernet, StarLAN, and PC Network. Note that we've put in here                */
/* what we actually see in the real world, not what IEEE would like              */
/*                                                                               */
manuf "VisTec" = ØØØØ22 /* Visual Technology, Inc.                               */
manuf "NwkGnl" = ØØØØ65 /* Network General Corp.                                 */
manuf "Prteon" = ØØØØ93 /* Proteon (bit-reversed from token ring!)               */
manuf "Amrstr" = ØØØØ9f /* Ameristar Technology                                  */
manuf "Wllflt" = ØØØØa2 /* Wellfleet                                             */
manuf "NCD "   = ØØØØa7 /* Network Computing Devices, Inc.                       */
manuf "NSC "   = ØØØØa9 /* Network Systems Corp.                                 */
manuf "RND "   = ØØØØbØ /* RAD Network Devices Ltd.                              */
manuf "Cimlin" = ØØØØb3 /* CIMlinc                                               */
manuf "WstDig" = ØØØØcØ /* Western Digital                                       */
manuf "HP EON" = ØØØØc6 /* H-P Intlgnt Networks Oper (EON)                       */
manuf "IBM "   = 1ØØØ5a /* (not bit-reversed from token ring)                    */
manuf "Intrln" = Ø2Ø7Ø1 /* Interlan, Inc.                                        */
manuf "NSC "   = Ø8ØØ17 /* Network System Corp.                                  */
manuf "Intrgr" = Ø8ØØ36 /* Intergraph                                            */
manuf "Univtn" = Ø8ØØ49 /* Univation                                             */
manuf "IBM "   = Ø8ØØ5a /* (bit-reversed from token ring)                        */
manuf "ComDes" = Ø8ØØ67 /* ComDesign                                             */


manuf "Xerox " = ØØØØaa  |  manuf "CMC "   = Ø2cf1f  |  manuf "DEC "   = Ø8ØØ2b
manuf "Dove "  = ØØØØb7  |  manuf "Bridge" = Ø8ØØØ2  |  manuf "Mtaphr" = Ø8ØØ2e
manuf "MIPS "  = ØØØØ6b  |  manuf "ACC "   = Ø8ØØØ3  |  manuf "Spider" = Ø8ØØ39
manuf "Ardent" = ØØØØ7a  |  manuf "Symblx" = Ø8ØØØ5  |  manuf "DCA "   = Ø8ØØ41
manuf "Cayman" = ØØØØ89  |  manuf "Apple " = Ø8ØØØ7  |  manuf "Sequnt" = Ø8ØØ47
manuf "TRW "   = ØØØØ2a  |  manuf "BBN "   = Ø8ØØØ8  |  manuf "Encore" = Ø8ØØ4c
manuf "Cisco " = ØØØØØc  |  manuf "H-P "   = Ø8ØØØ9  |  manuf "BICC "  = Ø8ØØ4e
manuf "NeXT "  = ØØØØØf  |  manuf "Nestar" = Ø8ØØØa  |  manuf "Ridge " = Ø8ØØ68
manuf "Sytek " = ØØØØ1Ø  |  manuf "Unisys" = Ø8ØØØb  |  manuf "SilGrf" = Ø8ØØ69
manuf "Novell" = ØØØØ1b  |  manuf "AT&T "  = Ø8ØØ1Ø  |  manuf "AT&T "  = Ø8ØØ6a
manuf "Altos " = ØØØØc8  |  manuf "Tktrnx" = Ø8ØØ11  |  manuf "Exceln" = Ø8ØØ6e
manuf "Gould " = ØØØØdd  |  manuf "Exceln" = Ø8ØØ14  |  manuf "Vtalnk" = Ø8ØØ7c
manuf "Acer "  = ØØØØe2  |  manuf "DG "    = Ø8ØØ1a  |  manuf "Xyplex" = Ø8ØØ87
manuf "Alantc" = ØØØØef  |  manuf "DG "    = Ø8ØØ1b  |  manuf "Kinetx" = Ø8ØØ89
manuf "Agilis" = ØØ8Ø5c  |  manuf "Apollo" = Ø8ØØ1e  |  manuf "Pyramd" = Ø8ØØ8b
manuf "Intel " = ØØaaØØ  |  manuf "Sun "   = Ø8ØØ2Ø  |  manuf "Xyvisn" = Ø8ØØ8d
manuf "U-B "   = ØØddØØ  |  manuf "NBI "   = Ø8ØØ22  |  manuf "Retix " = Ø8ØØ9Ø
manuf "U-B "   = ØØddØ1  |  manuf "CDC "   = Ø8ØØ25  |  manuf "DEC "   = aaØØØ3
manuf "3Com "  = Ø26Ø8c  |  manuf "TI "    = Ø8ØØ28  |  manuf "DECnet" = aaØØØ4
```

Figure 6–9. Manufacturer ID translation.

Manufacturer IDs in the table are shown as they appear in the computer once they have been captured. The IEEE assignment of IDs specifies the sequence in which bits are transmitted on the wire. For networks that transmit each byte low-order-bit first (such as Ethernet) the address you see after it has been captured is the byte-by-byte reverse of what was transmitted on the wire. For example, the code used by IBM, is transmitted on the network by the sequence 00010000 00000000 01011010. It appears to a token ring analysis server as 10 00 5A, but to an Ethernet server as 08 00 5A. See the discussion of the bit-reverse option in Chapter 5, page 5–27.

Data Exported to Printer or to File

The contents of the capture buffer (before or after filtering) may be sent to a file on the server or to a printer attached to the server or the console. If sent to a file, the contents may be in a printer format (with or without page titles and page numbers), or in the CSV format recognized by standard spread sheets. (For details of this procedure, see Chapter 5, "Exporting a Report on Frames in the Capture Buffer" on page 5–51 ff.)

The files thus produced have the extension .PRN for printer files and .CSV for spread sheet files.

Format of Saved Data Files

For some purposes of data analysis, you may want to upload data to the console for further analysis. After you've transferred data to a local computer, you can (for example) write a program that reads through a file of saved frames. The easiest format to work with is the ASCII file you get when you "print" the capture buffer to a file. Especially if you choose the option to omit page titles, the resulting file may contain the information you want in an easily-accessible format.

Alternatively, you may prefer to operate directly on the trace file that the Sniffer analyzer writes in response to the **Save data** command. This section describes the format of such a trace file.

Each trace file consists of sequences of variable length binary records. Since all 256 byte values are possible within the data, you cannot edit this file using an ordinary text editor.

The first 16 bytes of a trace file contain a text message identifying the file as one containing data collected by the Sniffer analyzer.[1] The message is followed by an end-of-file character (hex 1A, also called Ctrl-Z). Even if you accidentally type the file to the screen, or otherwise treat it as a text file, the display reaches a terminator before reaching unprintable characters.

Structures within the Data File

Following the text message string, the file contains an arbitrary number of variable-length records. Each record has a type, identified in its first two bytes. The three principal types are:

- Version record
- Frame record

[1] For historical reasons, the message in all such files is "TRSNIFF data", regardless of the network on which the frames were collected.

- End-of-file record

The first record in the file is always a version record, the last is always an end-of-file record, and those in between are usually (but not necessarily) frame records.

There is no explicit encoding of the file's total length (except as part of its directory entry).

Header

Every record, of any type, begins with the following header:

```
struct f_rec_struc {    /* Standard record header.                         */
int     type;           /* Type of this record. (Int = 2 bytes)            */
int     length;         /* Length of remainder of this record.             */
int     rsvd;           /* Reserved word, currently Ø.                     */
};
```

The header's first field indicates what type of record follows. The three principal types are identified by the following values:

```
#define REC_VERS        1   /* Version record (f_vers).                  */
#define REC_FRAME2      4   /* Frame data (f_frame2).                    */
#define REC_EOF         3   /* End-of-file record (no data follows).     */
```

Types other than these are reserved for future or other use. If you write a program to process data files, you should have it skip any record that isn't one of these types. The length field indicates how much data to skip.

Format of a Version Record

```
struct f_vers_struct    {
int       maj_vers;          /* Major version of the analyzer               */
int       min_vers;          /* Minor version of the analyzer               */
struct    date_struct date   /* Date & time (4 bytes, DOS format)           */
char      type;              /* What type of records follow.                */
char      network;           /* An indicator of the network type.           */
char      format;            /* An indicator of the format version.         */
char      timeunit;          /* An indicator of the frame timestamp unit.   */
int       rsvd[3];           /* Reserved words.                             */
};
```

The possible values of **network** are as follows:

```
#define     NETWORK_TRING       Ø   /* Token ring                          */
#define     NETWORK_ENET        1   /* Ethernet                            */
#define     NETWORK_ARCNET      2   /* ARCNET                              */
#define     NETWORK_STARLAN     3   /* StarLAN                             */
```

```
#define     NETWORK_PCNW          4    /* PC Network broadband      */
#define     NETWORK_LOCALTALK     5    /* LocalTalk                 */
#define     NETWORK_ZNET          6    /* Znet                      */
#define     NETWORK_SYNCHRO       7    /* WAN/Synchronous           */
```

The possible values of **timeunit** are as follows:

```
#define  TIMEUNIT_UNSPEC   0  /*  Unspecified; default by network type.  */
#define  TIMEUNIT_PC       1  /*   0.838096       microsecond units      */
#define  TIMEUNIT_3COM     2  /*  15.000000       microsecond units      */
#define  TIMEUNIT_MICOM    3  /*   0.500000       microsecond units      */
#define  TIMEUNIT_SYTEK    4  /*   2.000000       microsecond units      */
```

Format of a Frame Data Record

Each record starts with a header, as described above, followed by data in the following structure:

```
struct    f_frame2_struct    {
unsigned  time_low;         /* Low time, network-dependent units.        */
unsigned  time_mid;         /* Mid time, network-dependent units.        */
char      time_high;        /* High time, network-dependent units.       */
char      time_day;         /* Time in days since start of capture.      */
int       size;             /* Number of bytes actually written in this file
                               (may be less than frame's original length).  */
char      fs;               /* Frame error status bits.                  */
char      flags;            /* Buffer flags; for internal use.           */
int       true_size;        /* If nonzero, the size of the original frame
                                  (since this frame has been truncated).   */
int       rsvd;    }        /* Reserved; currently 0.
                               The frame data follows.                    */
```

All multibyte arithmetic fields (computed by the Sniffer analyzer during capture) are stored with the least significant byte first. Frame data are stored in the byte order transmitted.

Format of an End-of-file Record

The end-of-file record has no data; it consists only of the record header.

CHAPTER SEVEN: PROTOCOL INTERPRETER COMBINATIONS 7

Network General

Chapter 7. Protocol Interpreter Combinations

Chapter Overview

The Configuration Utility permits you to build or remove alternate versions of the Sniffer analyzer software on the server.

When you build a new analyzer, you specify the particular combination of protocol interpreter suites it will include. This permits you to create an analyzer that contains just the protocol interpreter suites you specify. Protocol suites are so numerous, and the individual protocol suites so large, that it is not possible to put all of them in a single analyzer.

The Configuration Utility automatically adds each new analyzer to the server's main selection menu, and automatically removes it from the menu when you delete one.

Need for Particular Combinations of Interpreters

With the Sniffer analysis server, Network General supplies *all* its protocol interpreter suites. However, you wouldn't really want all of them installed at once in the same analyzer. Moreover, they won't all fit in a single analyzer, anyway. So Network General delivers two versions of the analyzer and also gives you the ability to generate other combinations as needed.

The **Configure Server** menu includes an option for **Protocol Interpreter Combinations.** By selecting it, you can build whatever combinations you want, subject only to limitations on the total size of the resulting program.

Telling the Various Analyzers Apart

Looking at the left side of the server's **Main selection menu**, you'll always see an entry for the Sniffer monitor and entries for at least two versions of the Sniffer analyzer (Figure 7–1). Here's how you can tell the various analyzers apart. When you move the highlight to one of them, the *help* line below the panel lists its protocol interpreter suites. No two analyzers have the same list of interpreter suites. In Figure 7–1, the highlighted analyzer has suites for DECnet and X Windows.

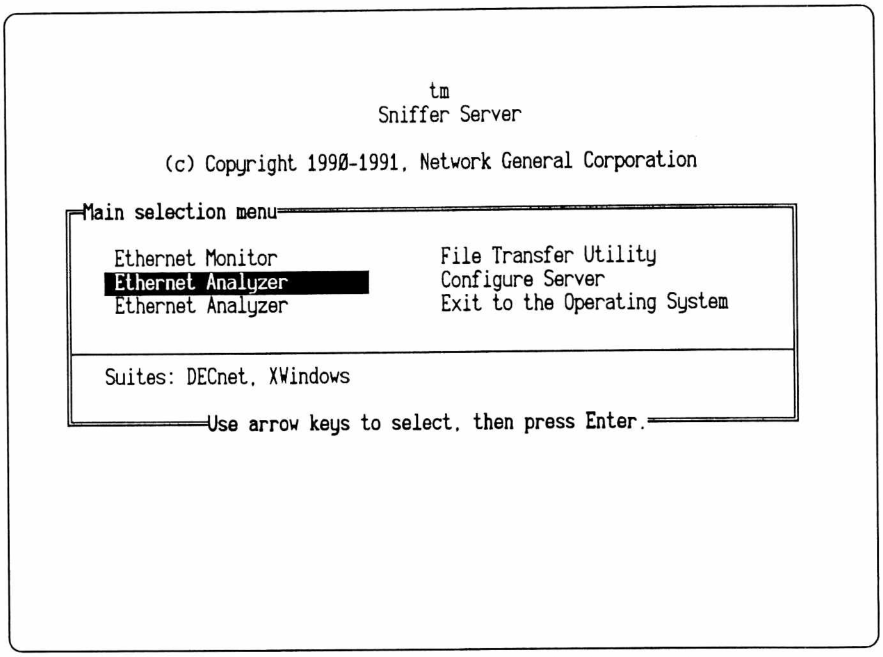

```
                                     tm
                               Sniffer Server

              (c) Copyright 1990-1991, Network General Corporation

 Main selection menu
   Ethernet Monitor                  File Transfer Utility
   Ethernet Analyzer                 Configure Server
   Ethernet Analyzer                 Exit to the Operating System

   Suites: DECnet, XWindows

            Use arrow keys to select, then press Enter.
```

Figure 7–1. Distinguishing between analyzers in the selection menu.

Looking at the right side of the server's **Main selection menu**, three choices are visible:

- File Transfer Utility (described in *Distributed Sniffer System: Installation and Operations Manual*)
- Configure Server (mostly described in *Distributed Sniffer System: Installation and Operations Manual*)
- Exit to the Operating System.

The second of these, **Configure server**, contains (among other things) the option **Protocol Interpreter Combinations**, described in the rest of this chapter.

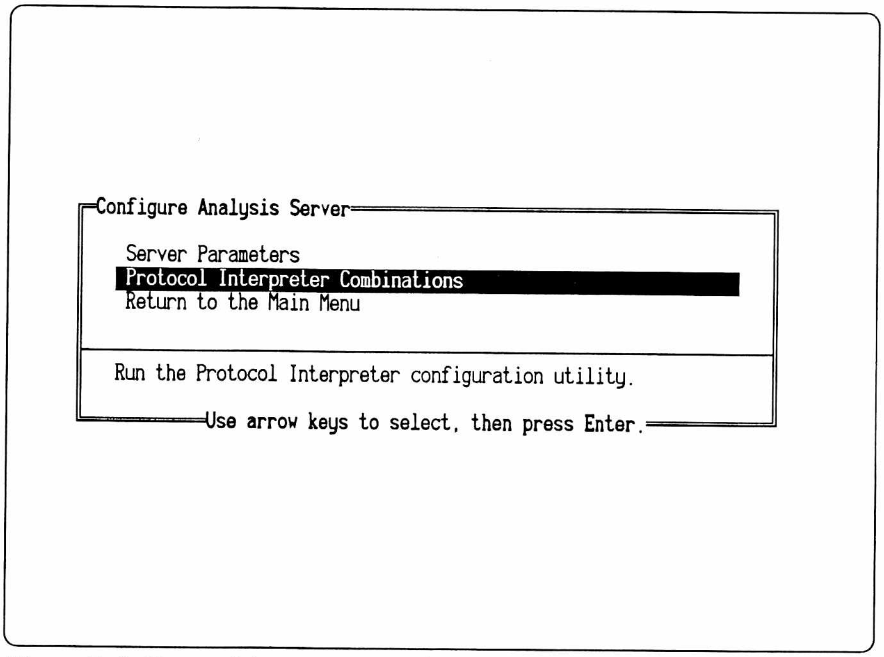

```
Configure Analysis Server

 Server Parameters
 Protocol Interpreter Combinations
 Return to the Main Menu

Run the Protocol Interpreter configuration utility.

         Use arrow keys to select, then press Enter.
```

Figure 7–2. Protocol interpreter combinations in the Configure menu.

You need Protocol Interpreter Combinations only to build a version of the Sniffer analyzer with a different combination of protocol interpreter suites. For example, suppose your server is initially configured with two analyzers, one with protocol interpreter suites for TCP/IP, Sun, ISO and X Windows, the other with DECnet and X Windows (Figure 7–3). Suppose you want in addition to build an analyzer having TCP/IP, DECnet and X.25.

As delivered by Network General		**As you wish to build**
Ethernet analyzer [file ENSN3BG0.EXE]	Ethernet analyzer [file ENSNSG80.EXE]	Ethernet analyzer [file ENSN2880.EXE]
1304 TCP/IP 1305 Sun 1307 DECnet 1309 VINES 1310 AppleTalk 1311 X-Windows	1301 IBM 1302 Novell 1303 XNS/MSNET 1306 ISO 1312 X.25	1304 TCP/IP 1307 DECnet 1312 X.25

Figure 7–3. Example: existing and desired combinations of protocols.

To build the new analyzer, you select **Protocol interpreter combinations** and tell the configuration program the protocols you want. It generates a new executable file. Like the other analyzers, the new one appears in the selection menu as **Ethernet analyzer**, but with its own distinct list of protocols when you highlight its name.[1]

The analysis server has a copy of the Microsoft Quick C Compiler. It resides in the subdirectory QC2 within the TOOLS directory. You don't run the compiler yourself; the configuration program does that automatically.

Building a New Analyzer

To build a Sniffer analyzer with a new combination of protocols

1. From the server's **Main selection menu**, move the highlight to **Configure server** and press Enter.

   Result: The server opens the submenu headed **Configure Analysis Server.**

2. Move the highlight to **Protocol Interpreter Combinations** and press Enter.

   Result: The server starts the Configuration Utility, and displays its initialization screen (Figure 7–4).

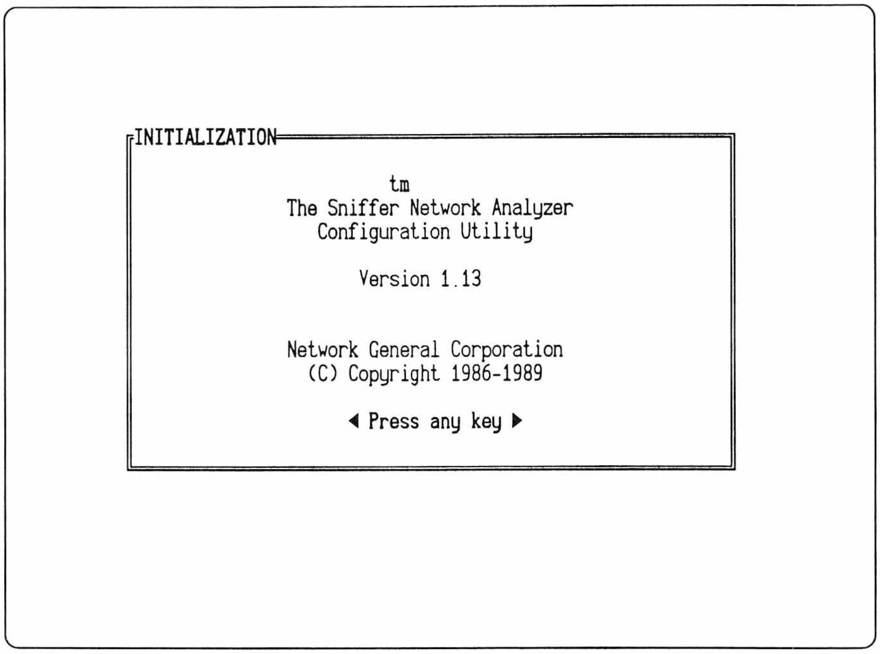

```
INITIALIZATION

                tm
   The Sniffer Network Analyzer
      Configuration Utility

           Version 1.13


    Network General Corporation
     (C) Copyright 1986-1989

         ◀ Press any key ▶
```

Figure 7–4. The Configuration Utility's initialization screen.

[1] If you exit to the operating system and examine the names of executable files, you'll see that each of the various analyzers appears to DOS to have a name in the form *xx*SN*xxxx*.EXE. Those names are explained in "Names for the Sniffer Analyzer's Executable Files" on page 6–5. They're also shown in square brackets in Figure 7–3

Result: When you acknowledge the Press any key message, the Configuration Utility brings up its main menu (Figure 7–5).

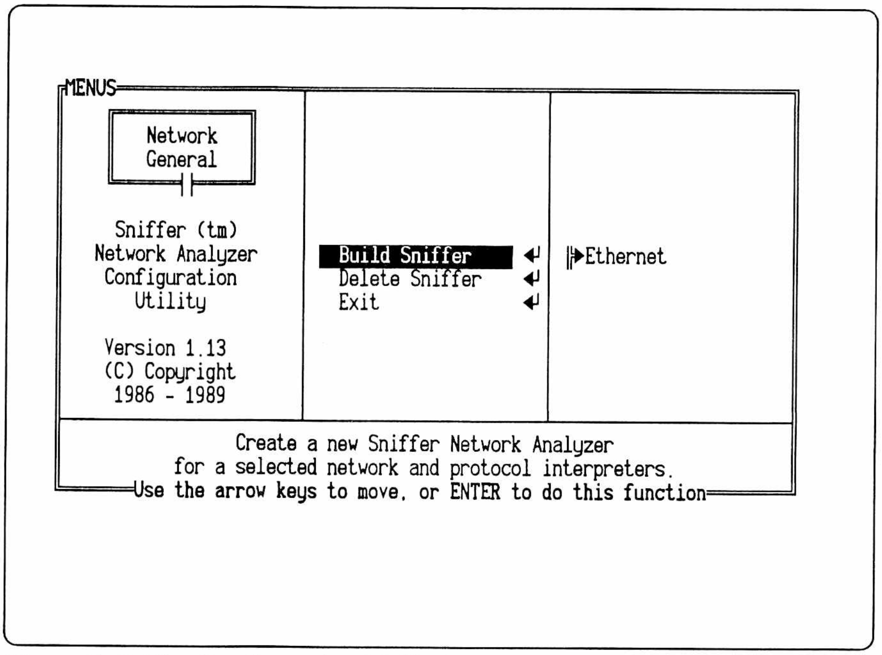

```
MENUS
  Network
  General

Sniffer (tm)
Network Analyzer      Build Sniffer  ↵    ▶Ethernet
Configuration         Delete Sniffer ↵
   Utility            Exit           ↵

Version 1.13
(C) Copyright
1986 - 1989

            Create a new Sniffer Network Analyzer
         for a selected network and protocol interpreters.
Use the arrow keys to move, or ENTER to do this function
```

Figure 7–5. The Configuration Utility's main menu.

3. In the Configuration Utility's main menu, move the highlight to **Build Sniffer**, and then to the right to select the target network.

   The next panel contains only one entry. It shows the type of network for which a new analyzer will be built (Ethernet, token ring, or WAN/synchronous). This may seem pointless: you have a radio control with only one choice. It's there because the same configuration program also runs on a portable Sniffer analyzer. A stand-alone Sniffer analyzer can have interface cards for several different networks. In that case, you'd get a real choice here.

4. From the network selection, move to the right to select the protocol interpreters you want (Figure 7–6).

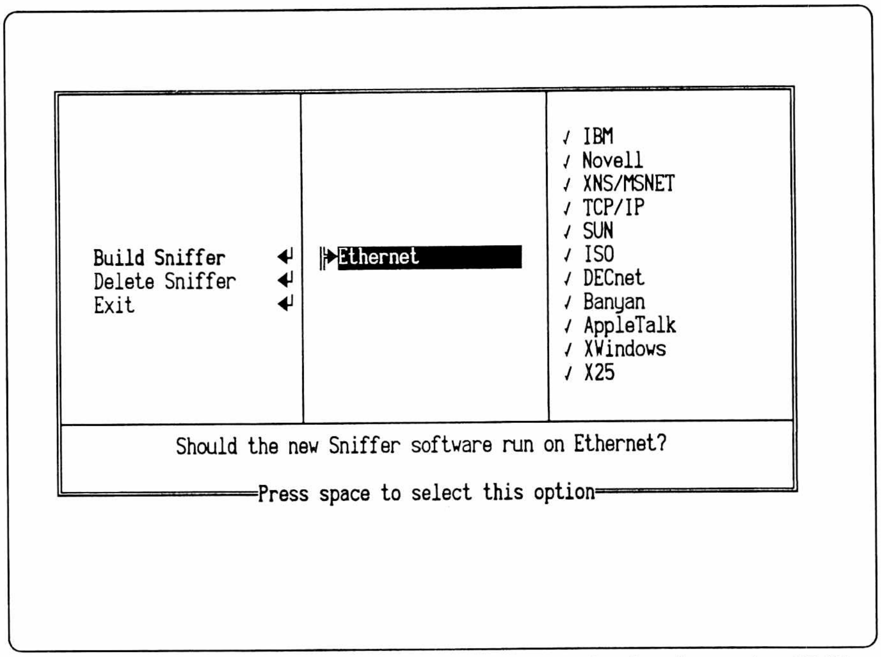

```
                                            √ IBM
                                            √ Novell
                                            √ XNS/MSNET
                                            √ TCP/IP
                                            √ SUN
Build Sniffer    ↵   ▶Ethernet              √ ISO
Delete Sniffer   ↵                          √ DECnet
Exit             ↵                          √ Banyan
                                            √ AppleTalk
                                            √ XWindows
                                            √ X25

           Should the new Sniffer software run on Ethernet?

                  Press space to select this option
```

Figure 7–6. Choices of network and interpreter in the Configuration Utility.

5. Within the list of protocol interpreters, move the highlight vertically to select a protocol interpreter and press Spacebar to toggle between √ (include) and X (exclude). Press Alt-Spacebar to reverse all the settings.

6. Make a judicious choice of protocols. It's hard to give specific advice here. Obviously, you want to include protocols likely to occur together. At the same time you don't want to include any that are unnecessary, since one of the objectives is to limit the total size of the analyzer to be built.

   The various protocol interpreter suites differ in size. The space required for a particular combination is not the sum of their separate parts. To be sure that what you request will fit into memory, you'll need to experiment. Even when the compilation runs successfully, the resulting executable may still be too large to execute. If so, you'll get an explanatory message when you try to start it.

7. Once you've selected the protocol suites to be included, move the highlight back to **Build Sniffer** and press Enter.

   Result: The Configuration Utility displays a description of the Sniffer analyzer you propose to build and asks you to confirm. (If you ask for a combination of protocols that already exists, it points out that your proposed analyzer will replace an existing one, and asks you to confirm.)

8. Confirm your wish to build a new analyzer by pressing Enter, or halt the build by pressing Esc.

Result: When you press Enter to proceed, the configuration program puts up a screen headed **Build progress**. As it starts each phase of the build, it reports its progress by adding a check mark beside the name of that stage. There are four stages:

- Building the configuration file
- Compiling
- Linking
- Post-link processing

At completion, the configuration program displays the message

```
xxx Sniffer Network Analyzer successfully built.
```

If the utility encounters an error during the process, you'll see a summary error message. This message may tell you where the Configuration Utility has written a more detailed error report.

When you exit from the Configuration Utility, you are returned to the selection menu. The Sniffer analyzer that you just built has now been added to the menu. It's ready to run in the same way as the other analyzers already installed.

Deleting an Analyzer

The menu in which you build a new analyzer also contains the choices to delete one. The task is simpler: you only have to identify the analyzer you no longer want.

To delete an existing Sniffer analyzer

1. From the server's **Main selection menu**, move the highlight to **Configure server** and press Enter.

   Result: The server opens the submenu headed **Configure Analysis Server.**

2. Move the highlight to **Protocol Interpreter Combinations** and press Enter.

   Result: The server starts the Configuration Utility, and displays its initialization screen (Figure 7–4).

When you acknowledge the `Press any key` message, the Configuration Utility brings up its main menu (Figure 7–5).

3. In the Configuration Utility's main menu, move the highlight to **Delete Sniffer**, and then to the right to select the analyzer to be deleted.

The next panel contains a list of the analyzers that currently exist on the server. The description of each analyzer shows first the network it monitors, and (in the lower panel) the protocol interpreter suites included in the highlighted analyzer. (That's why —if you look only at the center panel— the list of analyzers has several repetitions of the network name.)

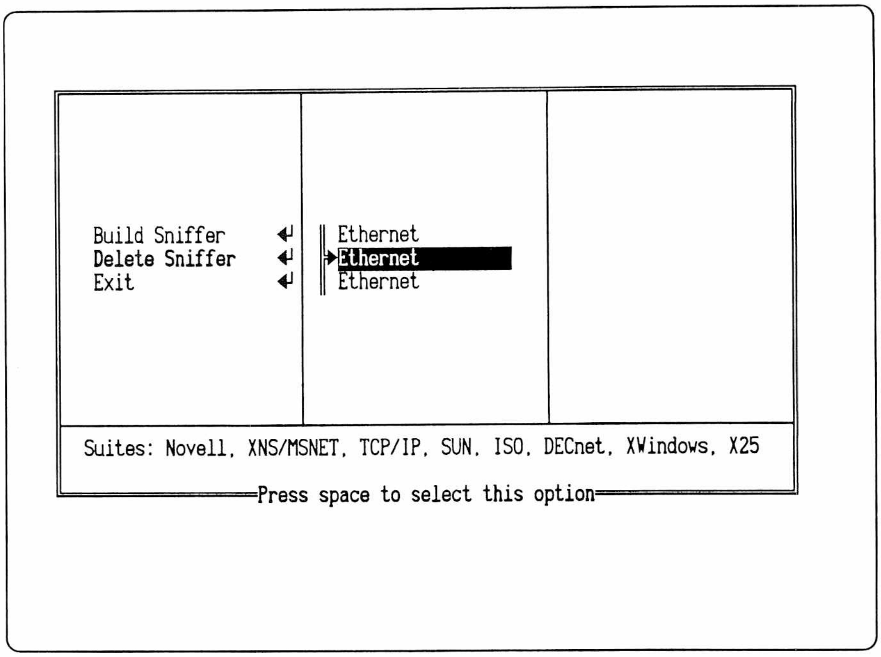

```
Suites: Novell, XNS/MSNET, TCP/IP, SUN, ISO, DECnet, XWindows, X25
═══════════Press space to select this option═══════════
```

Figure 7–7. Identifying an analyzer to be deleted.

4. Identify the analyzer to be deleted. As you move the highlight vertically from one analyzer to the next, the lower panel shows the list of protocol interpreters for the one you've highlighted. No two analyzers can have the same list of interpreters.

   To select the analyzer you want to delete, highlight its entry and press Spacebar to move the radio control's ║▶ to that line.

5. Having selected the analyzer to be deleted, return the highlight to **Delete Sniffer** and press Enter.

   Result: The configuration program displays a warning message (Figure 7–8). The message repeats the list of protocol interpreter suites of the analyzer in question. Press Enter to proceed with the deletion, or Esc to cancel. (If you delete an analyzer and subsequently decide you want it again, you have only to run the utility to rebuild it.)

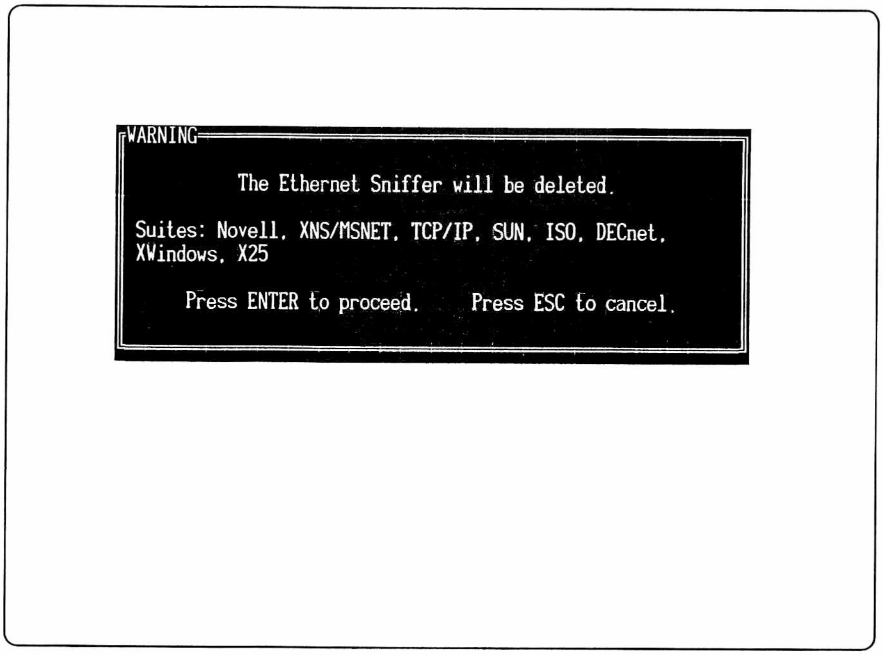

```
WARNING
            The Ethernet Sniffer will be deleted.

Suites: Novell, XNS/MSNET, TCP/IP, SUN, ISO, DECnet,
XWindows, X25

      Press ENTER to proceed.     Press ESC to cancel.
```

Figure 7–8. Warning before deletion of an analyzer.

DISTRIBUTED SNIFFER SYSTEM™

INDEX

Network General

Index

Network General

We solve network problems.™

Network General Corporation
4200 Bohannon Drive
Menlo Park, California 94025
(415) 688-2700

Network General Europe
Belgicastraat 4
1930 Zaventem, Belgium
32-2-725-6030